Recurrence of Nonlinear Control Systems: Entropy and Bit Rates

Hussein Sibai
Washington University in St. Louis

Enrique Mallada
Johns Hopkins University

Hybrid Systems: Computation and Control (HSCC)
CPS-IoT Week, Hong Kong
May 16, 2024

A World of Success Stories

2017 Google DeepMind’s DQN

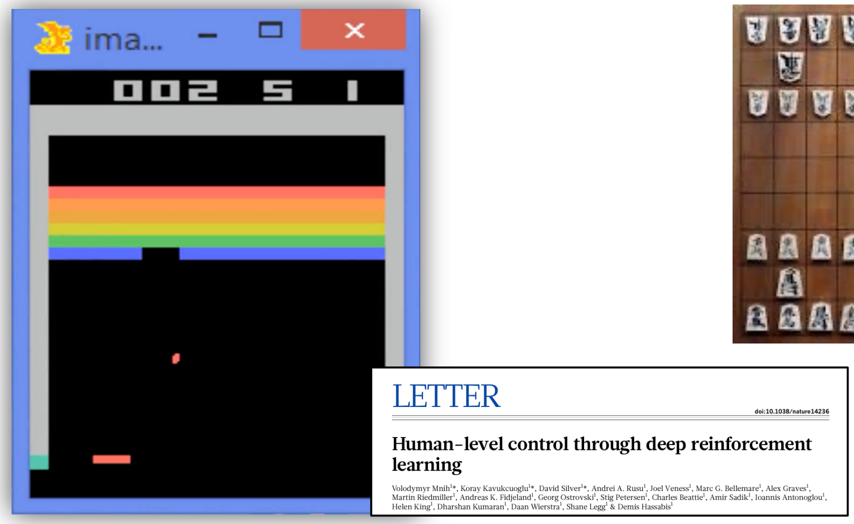

2017 AlphaZero – Chess, Shogi, Go

2019 AlphaStar – Starcraft II

Boston Dynamics

OpenAI – Rubik’s Cube

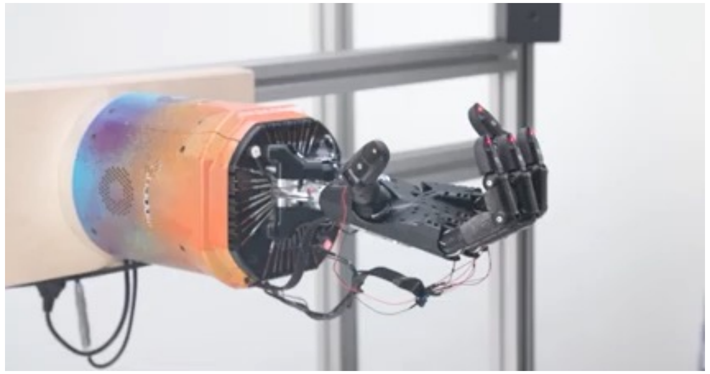

Waymo

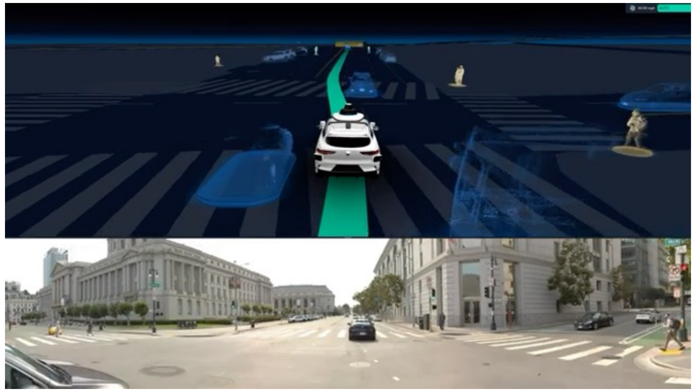

Reality Kicks In

Angry Residents, Abrupt Stops: Waymo Vehicles Are Still Causing Problems in Arizona

RAY STERN | MARCH 31, 2021 | 8:26AM

GARY MARCUS BUSINESS 08.14.2019 09:00 AM

DeepMind's Losses and the Future of Artificial Intelligence

Alphabet's DeepMind unit, conqueror of Go and other games, is losing lots of money. Continued deficits could imperil investments in AI.

AARIAN MARSHALL BUSINESS 12.07.2020 04:06 PM

Uber Gives Up on the Self-Driving Dream

The ride-hail giant invested more than $1 billion in autonomous vehicles. Now it's selling the unit to Aurora, which makes self-driving tech.

OpenAI disbands its robotics research team

Kyle Wiggers @Kyle_L_Wiggers July 16, 2021 11:24 AM

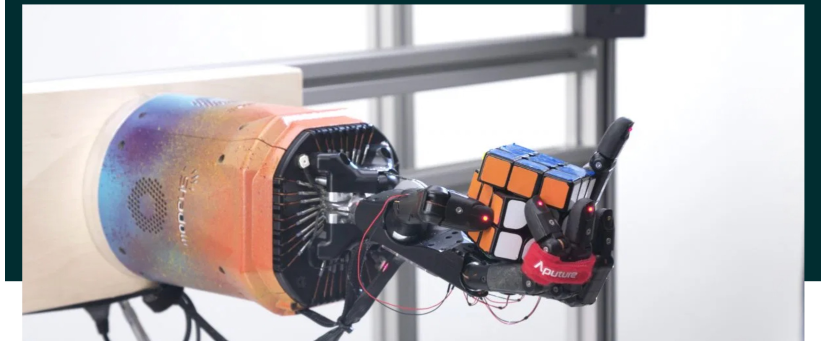

Self-driving Uber car that hit and killed woman did not recognize that pedestrians jaywalk

The automated car lacked "the capability to classify an object as a pedestrian unless that object was near a crosswalk," an NTSB report said.

Core challenge: The curse of dimensionality

- **Statistical: Sampling in d dimension with resolution ϵ**

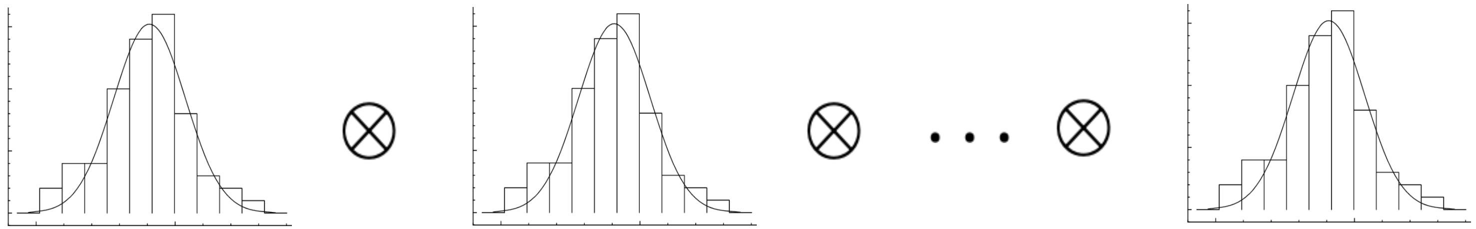

Sample complexity:

$$O(\varepsilon^{-d})$$

For $\epsilon = 0.1$ and $d = 100$, we would need 10^{100} points.

Atoms in the universe: 10^{78}

- **Computational: Verifying non-negativity of polynomials**

Copositive matrices:

$$[x_1^2 \dots x_d^2]A[x_1^2 \dots x_d^2]^{\mathrm{T}} \geq 0$$

Murty&Kadabi [1987]: Testing co-positivity is NP-Hard

Sum of Squares (SoS):

$$z(x)^T Q z(x) \geq 0, \quad z_i(x) \in \mathbb{R}[x], \ x \in \mathbb{R}^d, Q \succcurlyeq 0$$

Artin [1927] (Hilbert's 17th problem):

Non-negative polynomials are sum of square of *rational* functions

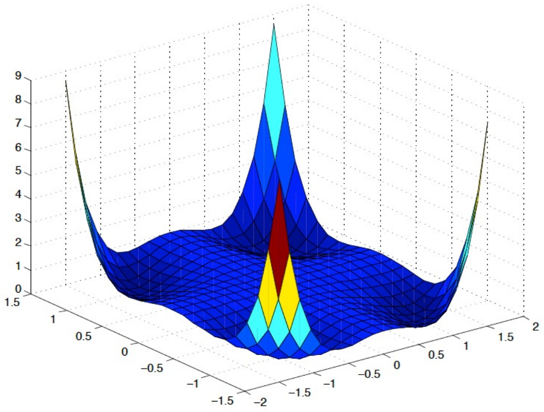

Motzkin [1967]:

$p = x^4y^2 + x^2y^4 + 1 - 3x^2y^2$

is nonnegative,

not a sum of squares,

but $(x^2 + y^2)^2 p$ is SoS

Question: Are we asking too much?

- Analysis tools build on a strict and exhaustive notion of ***invariance***

  Q: Can we substitute invariance with less restrictive notions?

  [arXiv '22] Shen, Bichuch, M - [CDC '23] Siegelmann, Shen, Paganini, M

- Certificates impose conditions on the entire duration of the trajectory

  Q: Can we provide guarantees based on only localized trajectory information?

  [arXiv '22] Shen, Bichuch, M - [CDC '23] Siegelmann, Shen, Paganini, M

- Control synthesis usually aims for the ***best*** (optimal) controller

  Q: Is there any gain in focusing on weaker requirements from the get-go?

  [HSCC 24] Sibai, M - - [CDC '23] Siegelmann, Shen, Paganini, M

[arXiv 22] Shen, Bichuch, M, *Model-free Learning of Regions of Attraction via Recurrent Sets,* **CDC 2022,** journal preprint arXiv:2204.10372.
[CDC 23] Siegelmann, Shen, Paganini, M, *A recurrence-based direct method for stability analysis and GPU-based verification of non-monotonic Lyapunov functions*, **CDC 2023**
[HSCC 24] Sibai, M, *Recurrence of nonlinear control systems: Entropy and bit rates,* **HSCC, 2024**

Question: Are we asking too much?

- Analysis tools build on a strict and exhaustive notion of ***invariance***
  Q: Can we substitute invariance with less restrictive notions?
  [arXiv ‘22] Shen, Bichuch, M - [CDC ’23] Siegelmann, Shen, Paganini, M

- Certificates impose conditions on the entire duration of the trajectory
  Q: Can we provide guarantees based on only localized trajectory information?
  [arXiv ‘22] Shen, Bichuch, M - [CDC ’23] Siegelmann, Shen, Paganini, M

- Control synthesis usually aims for the ***best*** (optimal) controller
  Q: Is there any gain in focusing on weaker requirements from the get-go?
  [HSCC 24] Sibai, M - [CDC ’23] Siegelmann, Shen, Paganini, M

[arXiv 22] Shen, Bichuch, M, *Model-free Learning of Regions of Attraction via Recurrent Sets,* **CDC 2022,** journal preprint arXiv:2204.10372.
[CDC 23] Siegelmann, Shen, Paganini, M, *A recurrence-based direct method for stability analysis and GPU-based verification of non-monotonic Lyapunov functions*, **CDC 2023**
[HSCC 24] Sibai, M, *Recurrence of nonlinear control systems: Entropy and bit rates*, HSCC, 2024

Outline

- Invariance: Merits and trade-offs
- Letting things go, and come back: Recurrent sets
- Analysis using recurrent sets
  - Approximating regions of attractions
  - Stability analysis via non-monotonic Lyapunov functions
- Recurrence in nonlinear control systems
  - Entropy and bit rates of control recurrent sets

Outline

- **Invariance: Merits and trade-offs**
- Letting things go, and come back: Recurrent sets
- Analysis using recurrent sets
  - Approximating regions of attractions
  - Stability analysis via non-monotonic Lyapunov functions
- Recurrence in nonlinear control systems
  - Entropy and bit rates of control recurrent sets

Problem setup

Continuous time dynamical system: $\dot{x}(t) = f(x(t))$

- Initial condition $x_0 = x(0)$, solution at time t: $\phi(t, x_0)$.

Ω-Limit Set $\Omega(f)$:

$$x \in \Omega(f) \iff \exists\, x_0, \{t_n\}_{n\geq 0}, \text{ s.t. } \lim_{n\to\infty} t_n = \infty \text{ and } \lim_{n\to\infty} \phi(t_n, x_0) = x$$

Types of Ω-limit set

equilibrium

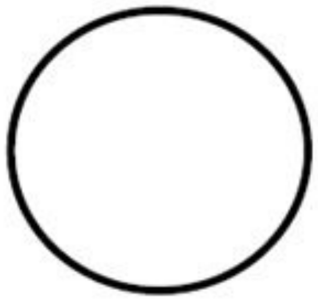

limit cycle

limit torus

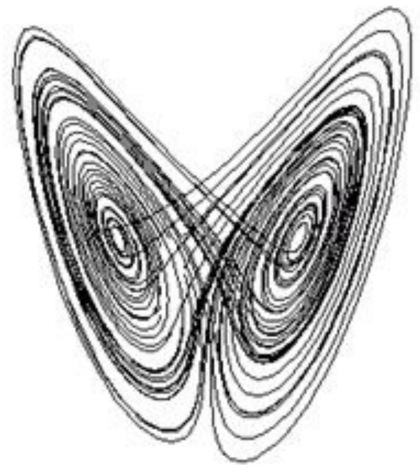

chaotic attractor

Problem setup

Continuous time dynamical system: $\dot{x}(t) = f(x(t))$

- Initial condition $x_0 = x(0)$, solution at time t: $\phi(t, x_0)$.

Ω-Limit Set $\Omega(f)$:

$$x \in \Omega(f) \iff \exists\, x_0, \{t_n\}_{n\geq 0}, \text{ s.t. } \lim_{n\to\infty} t_n = \infty \text{ and } \lim_{n\to\infty} \phi(t_n, x_0) = x$$

Illustrative Example

$$\begin{bmatrix}\dot{x}_1\\ \dot{x}_2\end{bmatrix} = \begin{bmatrix} x_2 \\ -x_1 + \frac{1}{3}x_1^3 - x_2\end{bmatrix}$$

$$\Omega(f) = \{(0,0), (-\sqrt{3}, 0), (\sqrt{3}, 0)\}$$ (equilibria)

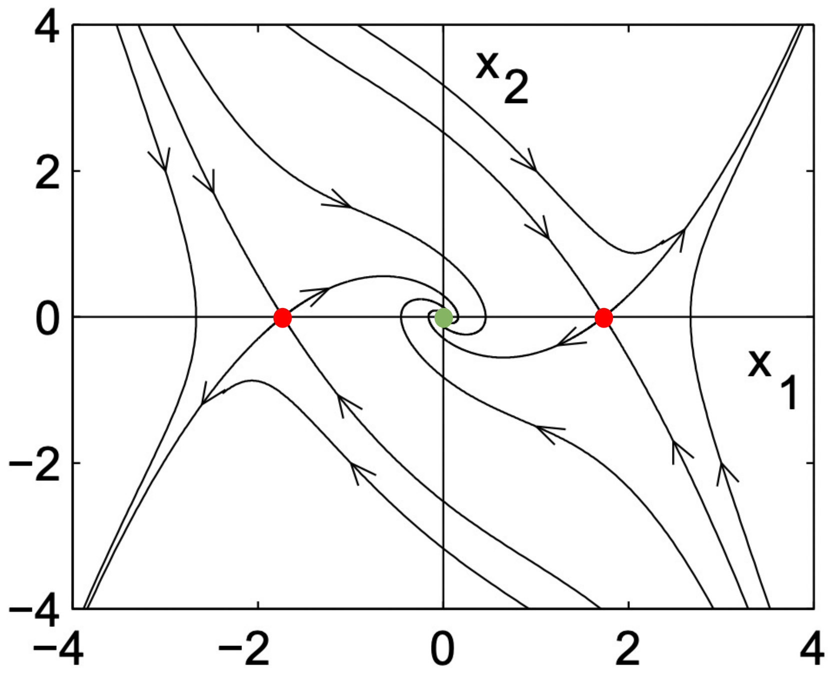

Problem setup

Continuous time dynamical system: $\dot{x}(t) = f(x(t))$

- Initial condition $x_0 = x(0)$, solution at time t: $\phi(t, x_0)$.
- The ω-limit set of the system: $\Omega(f)$

Region of attraction (ROA) of a set $S \subseteq \Omega(f)$:

$$\mathcal{A}(S) := \left\{ x \in \mathbb{R}^d \middle| \liminf_{t\to\infty} d(\phi(t,x), S) = 0 \right\}$$

Illustrative Example

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -x_1 + \frac{1}{3}x_1^3 - x_2 \end{bmatrix}$$

$$\Omega(f) = \{(0,0), (-\sqrt{3}, 0), (\sqrt{3}, 0)\}$$

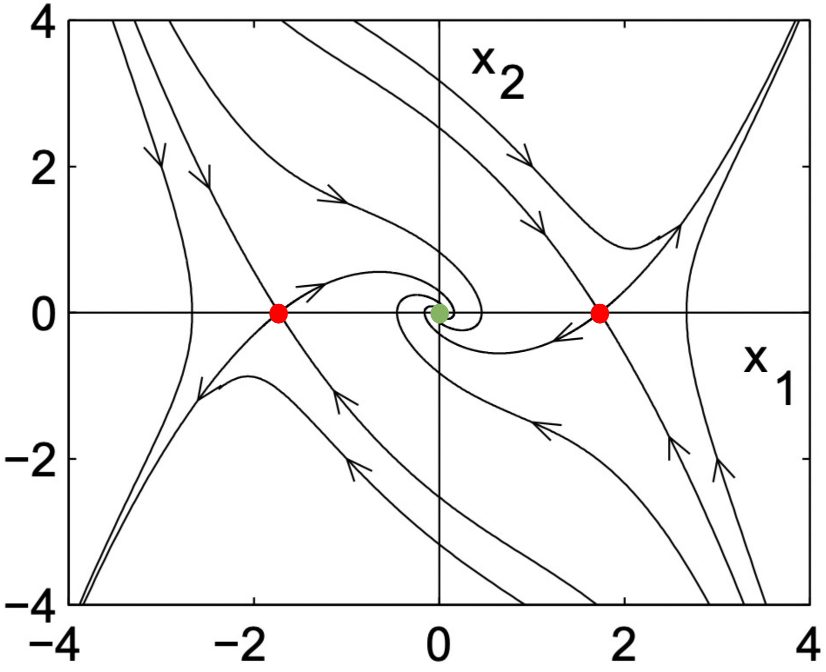

Problem setup

Continuous time dynamical system: $\dot{x}(t) = f(x(t))$

- Initial condition $x_0 = x(0)$, solution at time t: $\phi(t, x_0)$.
- The ω-limit set of the system: $\Omega(f)$

Region of attraction (ROA) of a set $S \subseteq \Omega(f)$:

$$\mathcal{A}(S) := \left\{ x \in \mathbb{R}^d \middle| \liminf_{t\to\infty} d(\phi(t,x), S) = 0 \right\}$$

Illustrative Example

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -x_1 + \frac{1}{3}x_1^3 - x_2 \end{bmatrix}$$

$$\Omega(f) = \{(0,0), (-\sqrt{3}, 0), (\sqrt{3}, 0)\}$$

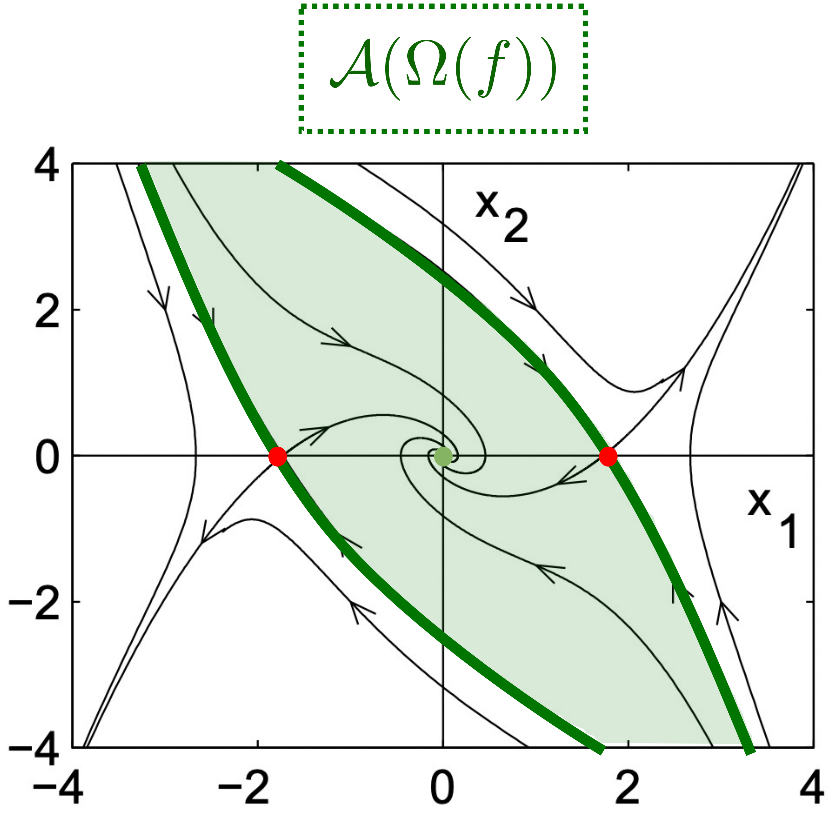

Problem setup

Continuous time dynamical system: $\dot{x}(t) = f(x(t))$

- Initial condition $x_0 = x(0)$, solution at time t: $\phi(t, x_0)$.
- The ω-limit set of the system: $\Omega(f)$

Region of attraction (ROA) of a set $S \subseteq \Omega(f)$:

$$\mathcal{A}(S) := \left\{ x \in \mathbb{R}^d \middle| \liminf_{t\to\infty} d(\phi(t,x), S) = 0 \right\}$$

Illustrative Example

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -x_1 + \frac{1}{3}x_1^3 - x_2 \end{bmatrix}$$

$$\Omega(f) = \{(0,0), (-\sqrt{3}, 0), (\sqrt{3}, 0)\}$$

Asymptotically stable equilibrium at $x^* = (0,0)$

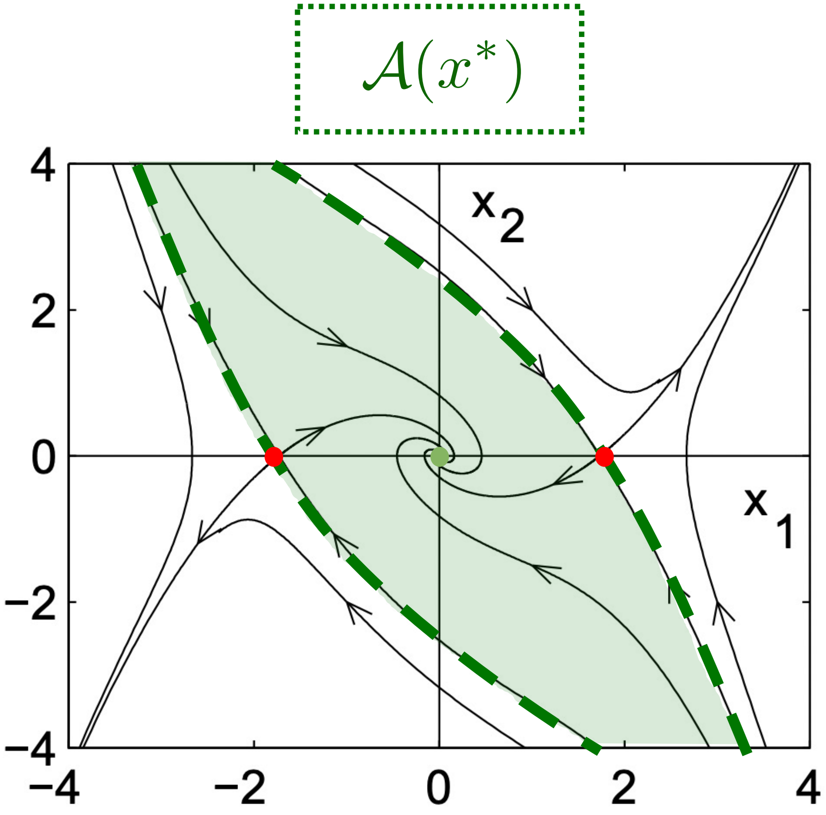

Problem setup

Continuous time dynamical system: $\dot{x}(t) = f(x(t))$

- Initial condition $x_0 = x(0)$, solution at time t: $\phi(t, x_0)$.
- The ω-limit set of the system: $\Omega(f)$

Region of attraction (ROA) of a set $S \subseteq \Omega(f)$:

$$\mathcal{A}(S) := \left\{ x \in \mathbb{R}^d \middle| \liminf_{t\to\infty} d(\phi(t,x), S) = 0 \right\}$$

Illustrative Example

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ -x_1 + \frac{1}{3}x_1^3 - x_2 \end{bmatrix}$$

$$\Omega(f) = \{(0,0), (-\sqrt{3},0), (\sqrt{3},0)\}$$

Unstable equilibria $\{(\sqrt{3},0), (-\sqrt{3},0)\}$

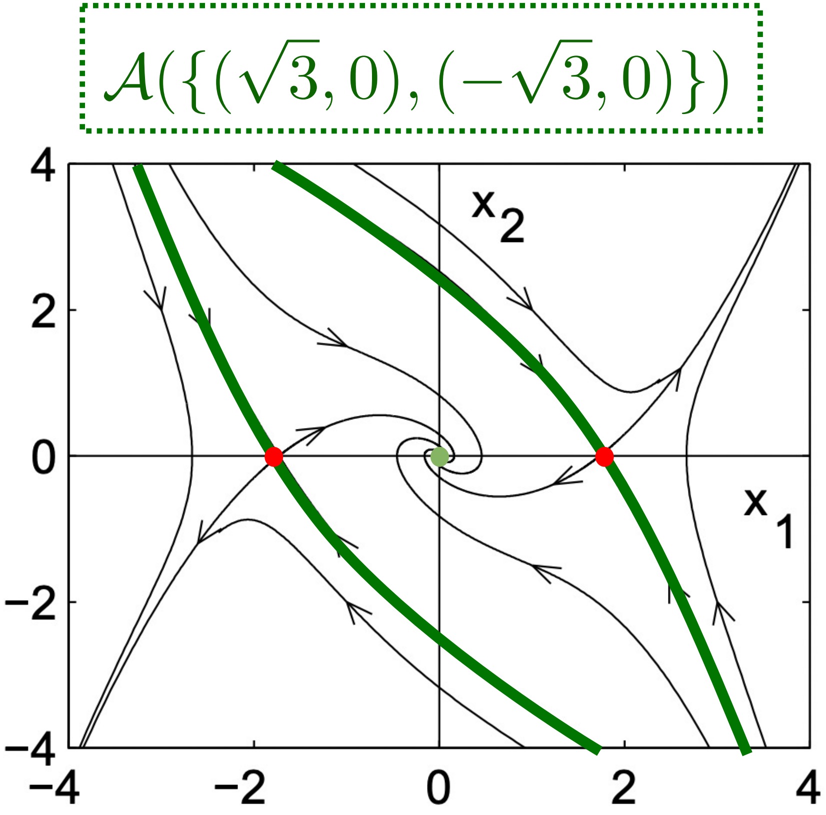

Invariant sets

A set $\mathcal{S} \subseteq \mathbb{R}^d$ is **positively invariant** if and only if: $x_0 \in \mathcal{S} \to \phi(t, x_0) \in \mathcal{S},\ \forall t \geq 0$

Any trajectory starting in the set remains in inside it for all times

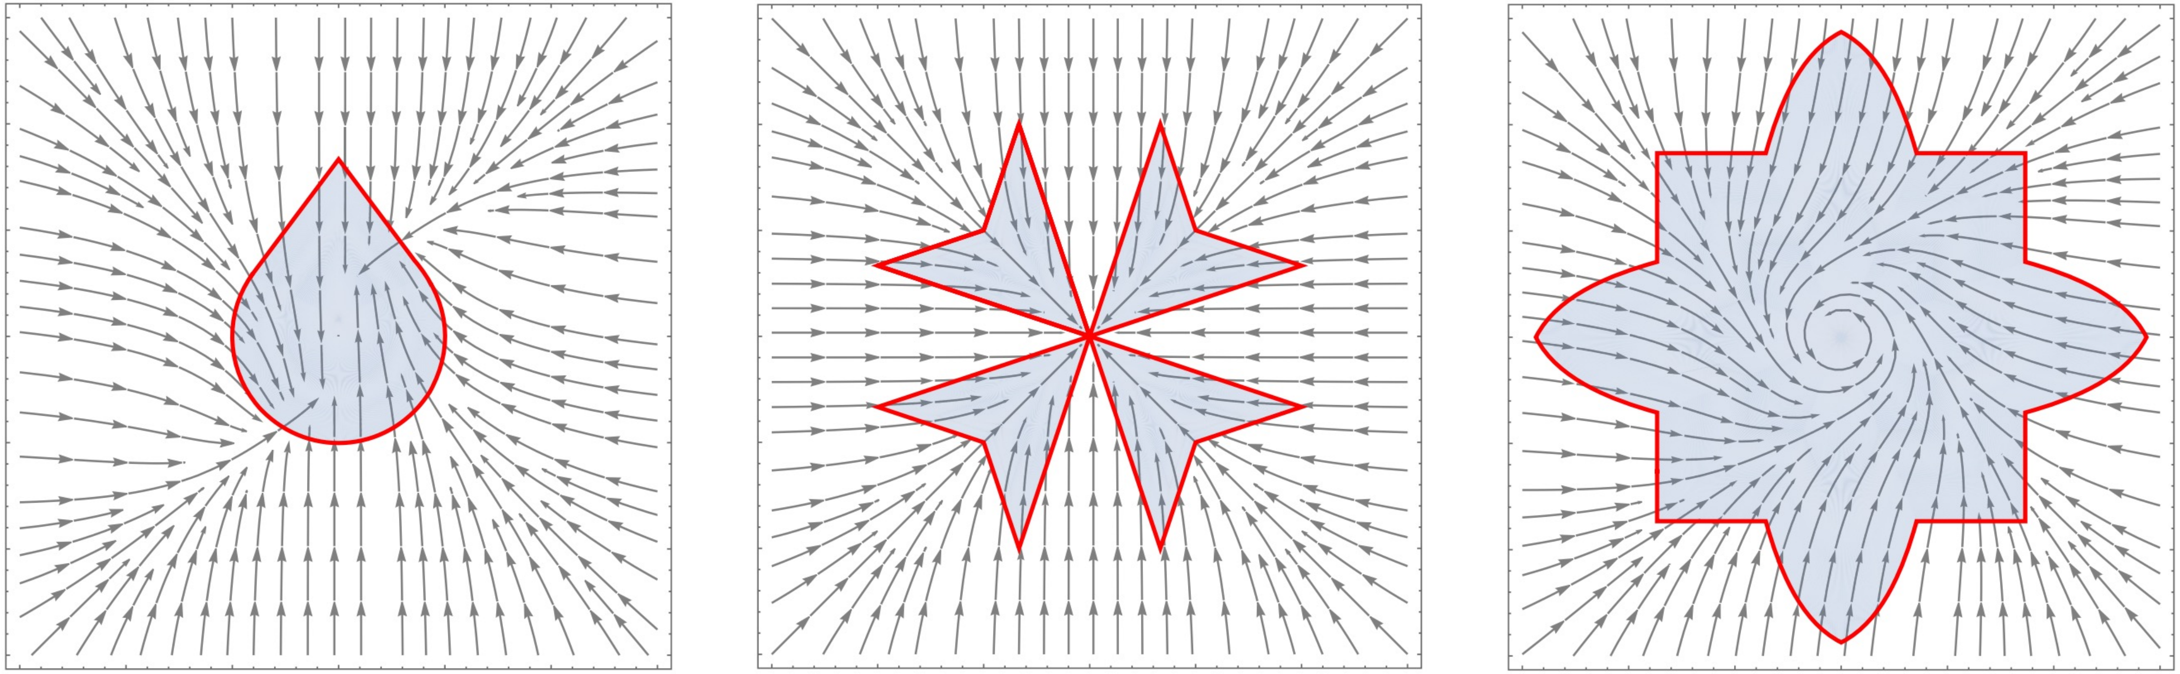

Source: K. Ghorbal, K. and A. Sogokon, Characterizing positively invariant sets: Inductive and topological methods. Journal of Symbolic Computation, 2022

Invariant sets: Merits

A set $\mathcal{S} \subseteq \mathbb{R}^d$ is **positively invariant** if and only if: $x_0 \in \mathcal{S} \to \phi(t, x_0) \in \mathcal{S},\ \forall t \geq 0$

Any trajectory starting in the set remains in inside it for all times

- **Invariant sets approximate regions of attraction**
  Compact invariant set $\mathcal{S}$ containing only $\{x^*\} = \Omega(f) \cap \mathcal{S}$ in the interior must be in the region of attraction $\mathcal{A}(x^*)$

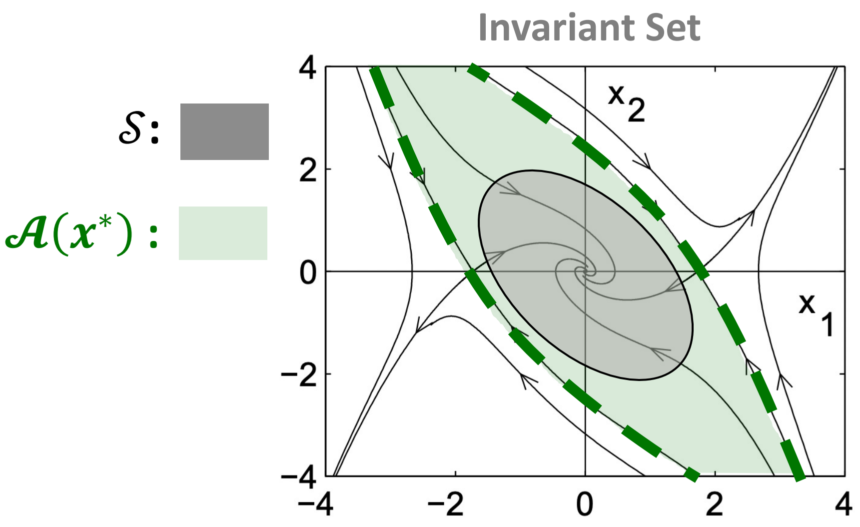

Invariant sets: Merits

A set $\mathcal{S} \subseteq \mathbb{R}^d$ is **positively invariant** if and only if: $x_0 \in \mathcal{S} \to \phi(t, x_0) \in \mathcal{S},\ \forall t \geq 0$

Any trajectory starting in the set remains in inside it for all times

- **Invariant sets approximate regions of attraction**
  Compact invariant set $\mathcal{S}$ containing only $\{x^*\} = \Omega(f) \cap \mathcal{S}$ in the interior must be in the region of attraction $\mathcal{A}(x^*)$

- **Invariant sets guarantee stability**
  Lyapunov stability: solutions starting "close enough" to the equilibrium (within a distance δ) remain "close enough" forever (within a distance ε)

- **Invariant sets further certify asymptotic stability via Lyapunov's direct method**
  Asymptotic stability: solutions that start close enough, remain close enough, and eventually converge to equilibrium.

Invariant Set

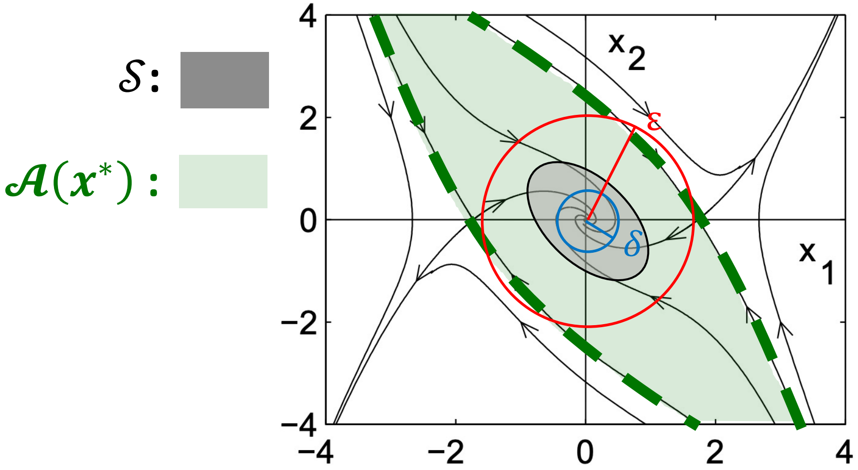

Lyapunov Functions

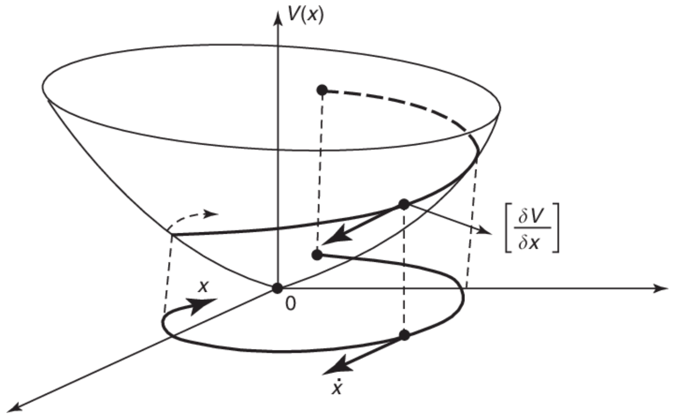

Invariant sets: Challenges

A set $\mathcal{S} \subseteq \mathbb{R}^d$ is **positively invariant** if and only if: $x_0 \in \mathcal{S} \rightarrow \phi(t, x_0) \in \mathcal{S},\ \forall t \geq 0$

Any trajectory starting in the set remains in inside it for all times

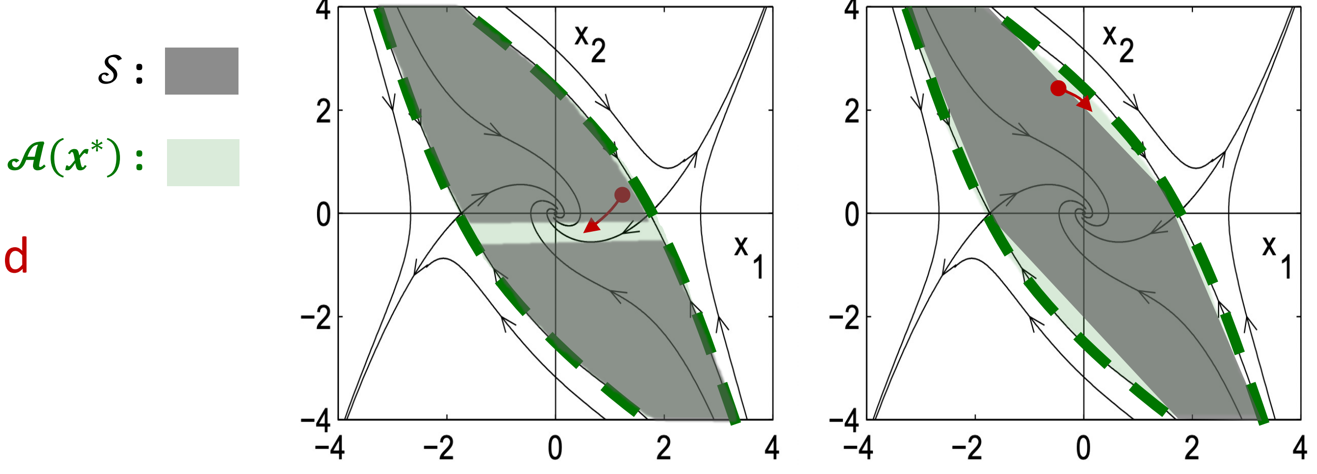

A not invariant trajectory:

- **$\mathcal{S}$ is topologically constrained**
  - If $\mathcal{S} \cap \Omega(f) = \{x^*\}$, then $\mathcal{S}$ is connected
- **$\mathcal{S}$ is geometrically constrained**
  - f should not point outwards for $x \in \partial\mathcal{S}$
- **$\mathcal{S}$ geometry can be wild**
  - $\mathcal{A}(\Omega(f))$ is not necessarily analytic!

Basin of $\mathcal{A}(\Omega(f))$

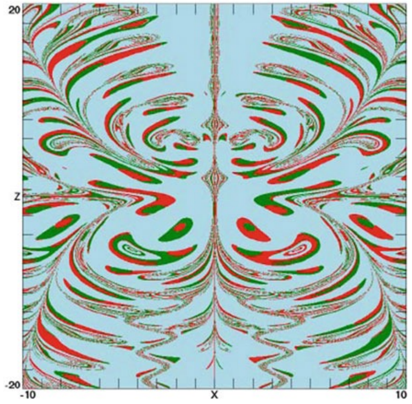

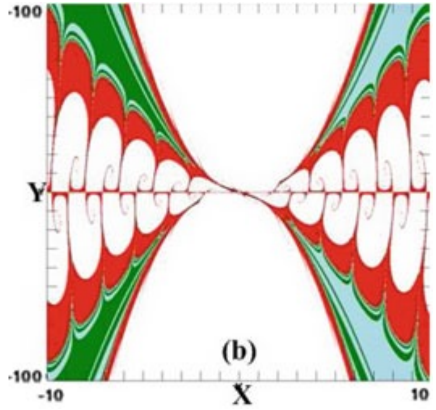

Outline

- Invariance: Merits and trade-offs
- **Letting things go, and come back: Recurrent sets**
- Analysis using recurrent sets
  - Approximating regions of attractions
  - Stability analysis via non-monotonic Lyapunov functions
- Recurrence in nonlinear control systems
  - Entropy and bit rates of control recurrent sets

Recurrent sets: Letting things go, and come back

A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **recurrent** if for any $x_0 \in \mathcal{R}$ and $t \geq 0$, $\exists t' \geq t$ s.t. $\phi(t', x_0) \in \mathcal{R}$.

Property of Recurrent Sets

- $\mathcal{R}$ need **not** be **connected**
- $\mathcal{R}$ does **not** require f to **point inwards** on all $\partial\mathcal{R}$

Recurrent sets, while not invariant, guarantee that solutions that start in this set, will come back **infinitely often, forever!**

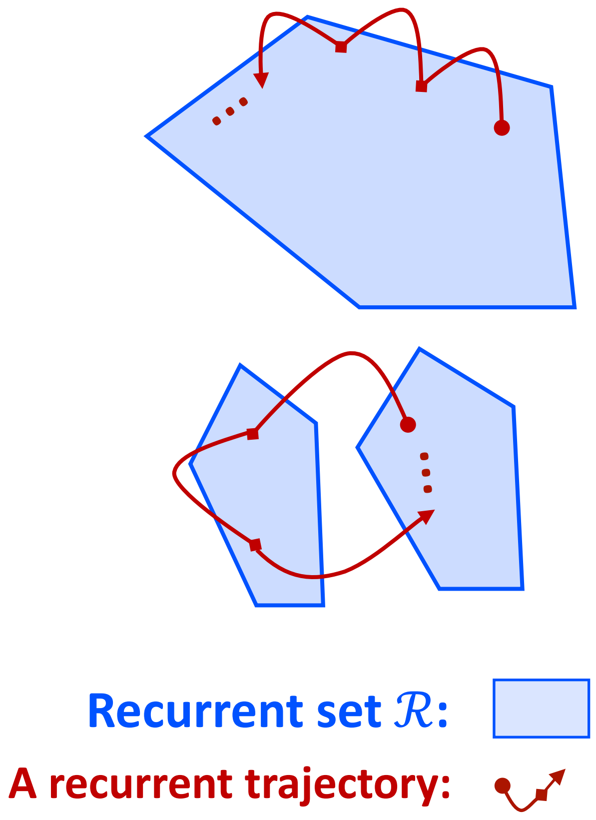

Recurrent sets: Letting things go, and come back

A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **recurrent** if for any $x_0 \in \mathcal{R}$ and $t \geq 0$, $\exists t' \geq t$ s.t. $\phi(t', x_0) \in \mathcal{R}$.

Previous two good inner approximations of $\mathcal{A}(x^*)$ are recurrent sets

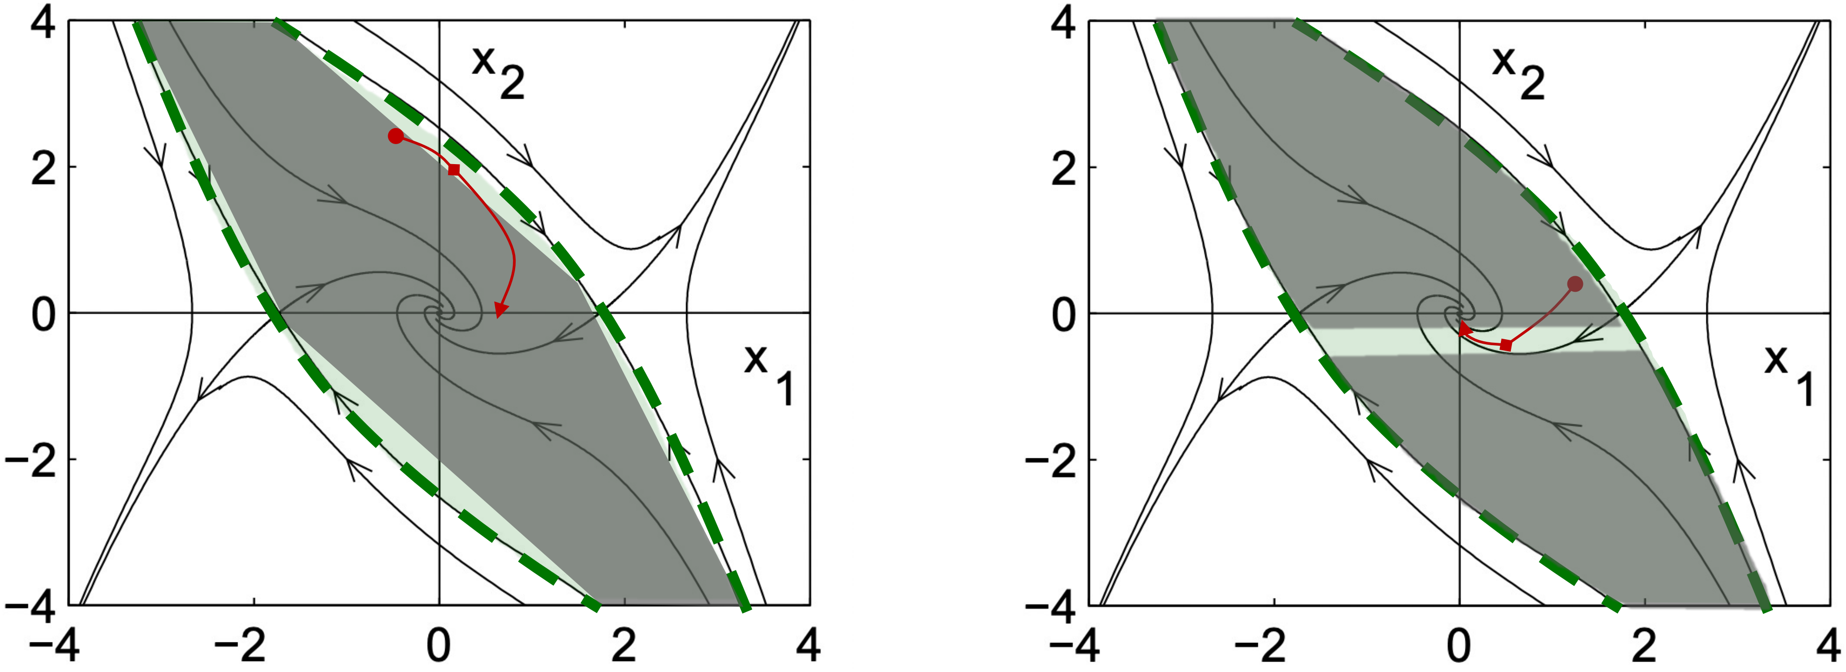

[arXiv 22] Shen, Bichuch, M, *Model-free Learning of Regions of Attraction via Recurrent Sets,* **CDC 2022,** journal preprint arXiv:2204.10372.

Recurrent sets: Letting things go, and come back

A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **recurrent** if for any $x_0 \in \mathcal{R}$ and $t \geq 0$, $\exists t' \geq t$ s.t. $\phi(t', x_0) \in \mathcal{R}$.

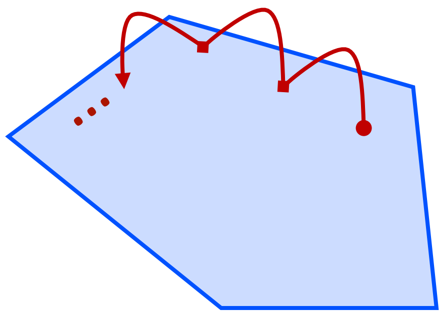

Property of Recurrent Sets

- $\mathcal{R}$ need **not** be **connected**
- $\mathcal{R}$ does **not** require f to **point inwards** on all $\partial\mathcal{R}$

> Recurrent sets, while not invariant, guarantee that solutions that start in this set, will come back **infinitely often, forever!**

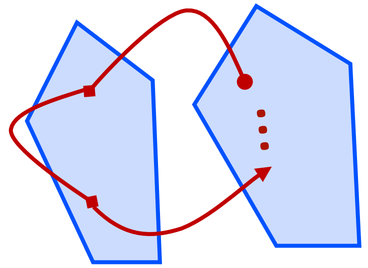

Question: Can we use recurrent sets as a substitute to invariant sets?

Outline

- Invariance: Merits and trade-offs
- Letting things go, and come back: Recurrent sets
- **Analysis using recurrent sets**
  - **Approximating regions of attractions**
  - Stability analysis via non-monotonic Lyapunov functions
- Recurrence in nonlinear control systems
  - Entropy and bit rates of control recurrent sets

Recurrent sets are subsets of the region of attraction

A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **recurrent** if for any $x_0 \in \mathcal{R}$ and $t \geq 0$, $\exists t' \geq t$ s.t. $\phi(t', x_0) \in \mathcal{R}$.

Theorem. *Let $\mathcal{R} \subset \mathbb{R}^d$ be a compact set satisfying $\partial\mathcal{R} \cap \Omega(f) = \emptyset$.*

Then:

$\mathcal{R}$ is *invariant* ➡ $\mathcal{R} \cap \Omega(f) \neq \emptyset$, $\mathcal{R} \subset \mathcal{A}(\mathcal{R} \cap \Omega(f))$

[arXiv 22] Shen, Bichuch, M, *Model-free Learning of Regions of Attraction via Recurrent Sets,* **CDC 2022,** journal preprint arXiv:2204.10372.

Recurrent sets are subsets of the region of attraction

A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **recurrent** if for any $x_0 \in \mathcal{R}$ and $t \geq 0$, $\exists t' \geq t$ s.t. $\phi(t', x_0) \in \mathcal{R}$.

Theorem. *Let $\mathcal{R} \subset \mathbb{R}^d$ be a compact set satisfying $\partial\mathcal{R} \cap \Omega(f) = \emptyset$.*

Then:

$\mathcal{R}$ *is recurrent* $\Longleftrightarrow$ $\mathcal{R} \cap \Omega(f) \neq \emptyset$, $\mathcal{R} \subset \mathcal{A}(\mathcal{R} \cap \Omega(f))$

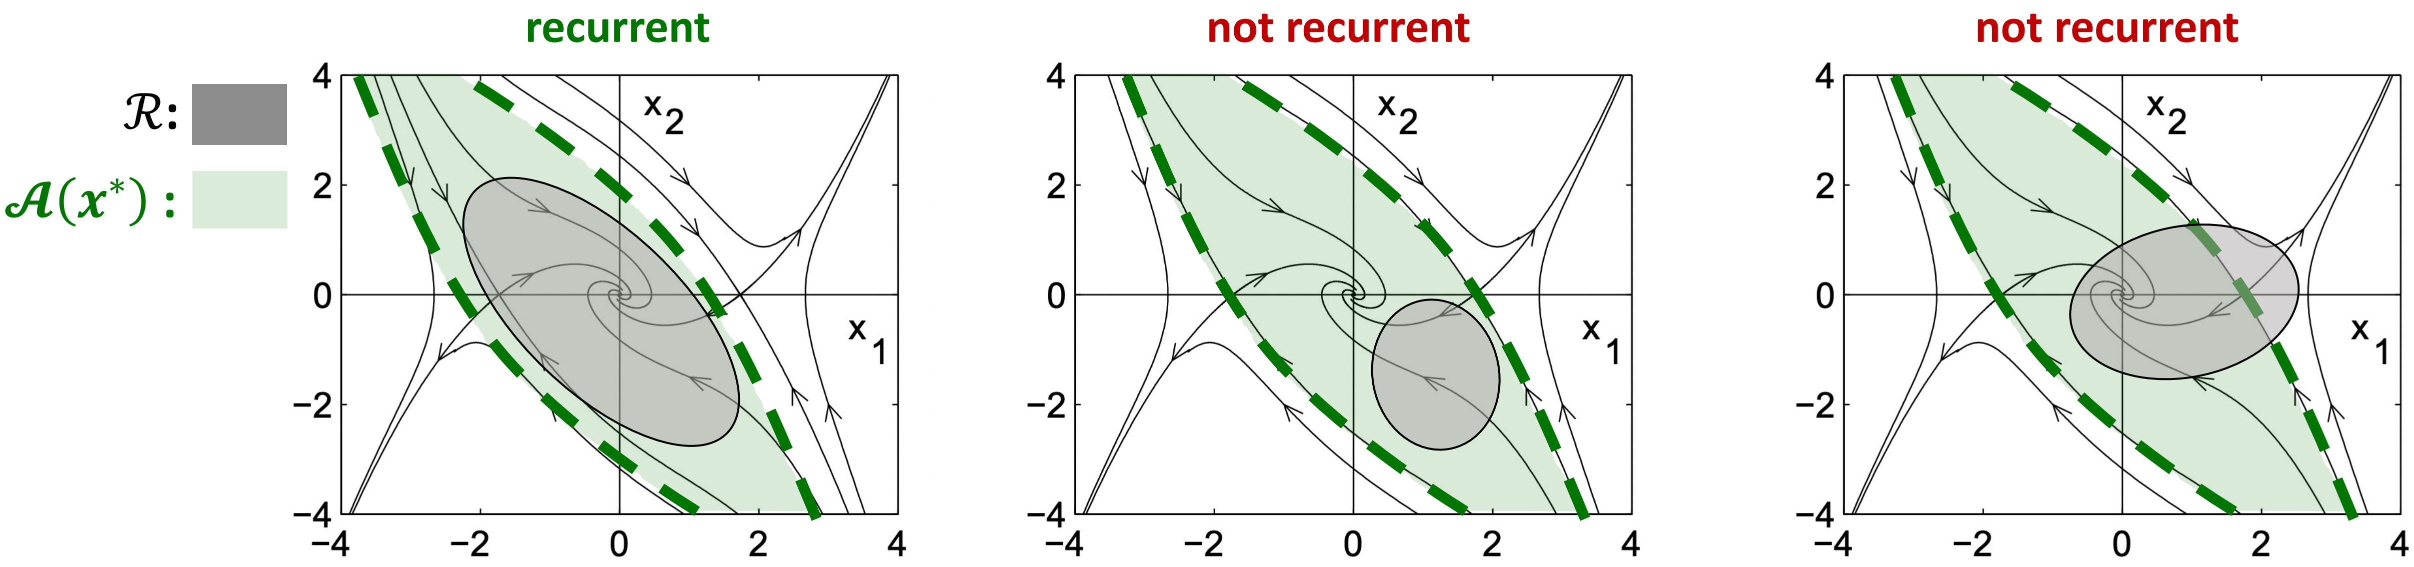

[arXiv 22] Shen, Bichuch, M, *Model-free Learning of Regions of Attraction via Recurrent Sets,* **CDC 2022,** journal preprint arXiv:2204.10372.

Learning Regions of Attractions via Recurrent Sets

Algorithm: Given h, k, and $\varepsilon > 0$:

- *Build approximation using unions of balls centered at* $x_1, \ldots, x_q$*, with* $x_1 = x^*$

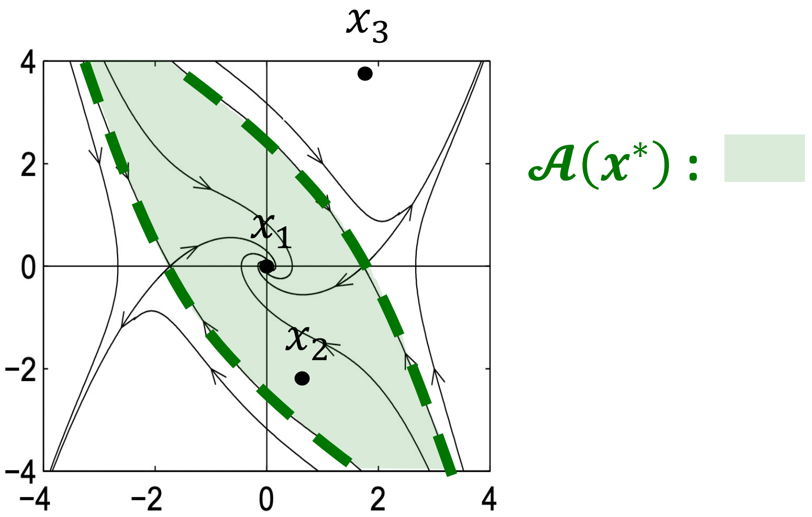

[arXiv 22] Shen, Bichuch, M, *Model-free Learning of Regions of Attraction via Recurrent Sets,* **CDC 2022,** journal preprint arXiv:2204.10372.

Learning Regions of Attractions via Recurrent Sets

Algorithm: Given h, k, and $\varepsilon > 0$:

- *Build approximation using unions of balls centered at* $x_1, \ldots, x_q$, *with* $x_1 = x^*$
- *Initial approximation:* $\mathcal{S}_0 = \cup_{q=1}^{h} \mathcal{S}_0^q$, *where* $\mathcal{S}_0^q = \left\{x \colon \left\|x - x_q\right\| \leq b_0^q\right\}$

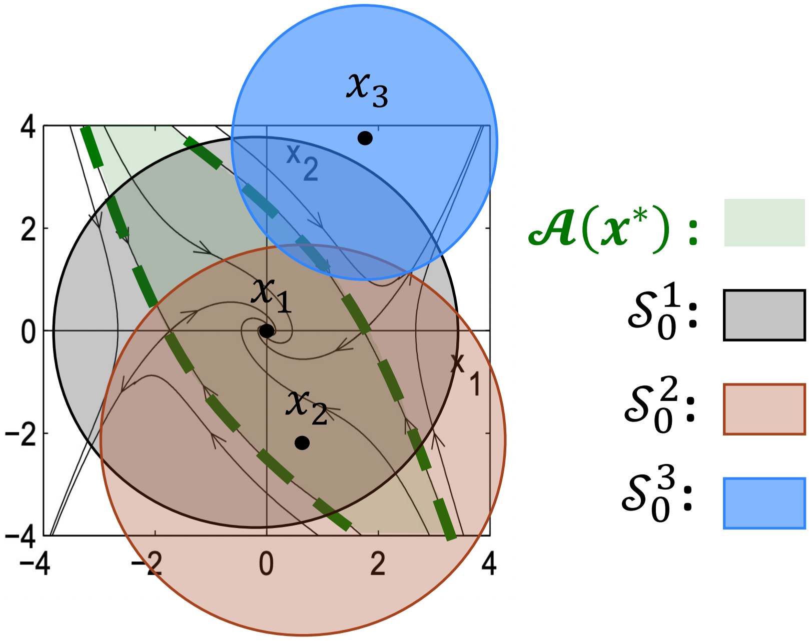

[arXiv 22] Shen, Bichuch, M, *Model-free Learning of Regions of Attraction via Recurrent Sets,* **CDC 2022,** journal preprint arXiv:2204.10372.

Learning Regions of Attractions via Recurrent Sets

Algorithm: Given h, k, and $\varepsilon > 0$:

- *Build approximation using unions of balls centered at* $x_1, \dots, x_q$, *with* $x_1 = x^*$
- *Initial approximation:* $\mathcal{S}_0 = \cup_{q=1}^{h} \mathcal{S}_0^q$, *where* $\mathcal{S}_0^q = \left\{x \colon \left\|x - x_q\right\| \leq b_0^q\right\}$

At each iteration l

- *Sample trajectories of duration* τ *from* $\mathcal{S}_l$ *until recurrence is violated (counter-example)*

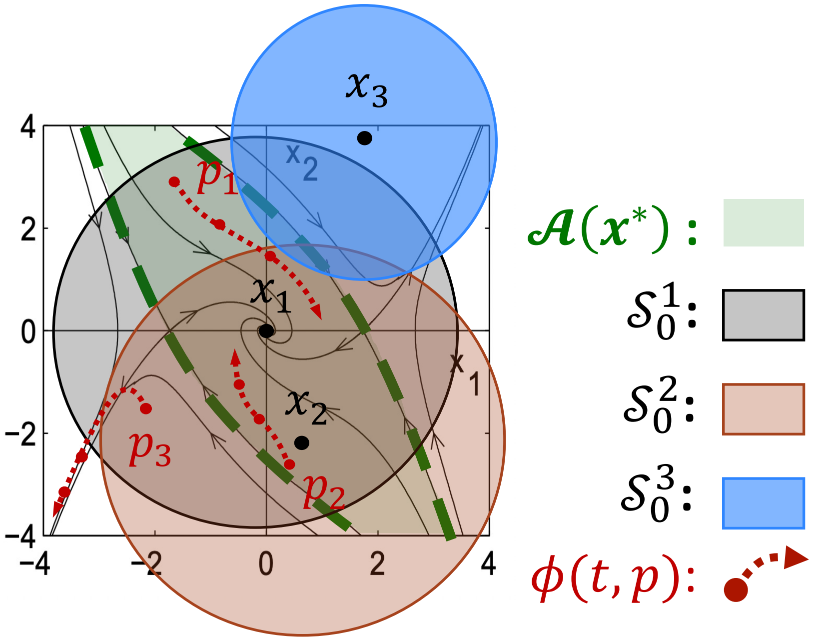

p_3: counter example

[arXiv 22] Shen, Bichuch, M, *Model-free Learning of Regions of Attraction via Recurrent Sets,* **CDC 2022,** journal preprint arXiv:2204.10372.

Learning Regions of Attractions via Recurrent Sets

Algorithm: Given h, k, and $\varepsilon > 0$:

- *Build approximation using unions of balls centered at* $x_1, \dots, x_q$, *with* $x_1 = x^*$
- *Initial approximation:* $\mathcal{S}_0 = \cup_{q=1}^{h} \mathcal{S}_0^q$, *where* $\mathcal{S}_0^q = \left\{x: \left\|x - x_q\right\| \leq b_0^q\right\}$

At each iteration l

- *Sample trajectories of duration* τ *from* $\mathcal{S}_l$ *until recurrence is violated (counter-example)*
- *Update approximation* $\mathcal{S}_{l+1}$ *to exclude counter-example neighborhood:* $p_j + \mathrm{B}_\varepsilon$

Sample complexity: $m \geq \frac{\mathrm{V}(\mathcal{S}_l + \mathrm{B}_\varepsilon)}{\mathrm{V}(\mathrm{B}_\varepsilon)} \log\left(\frac{1}{\delta}\right)$

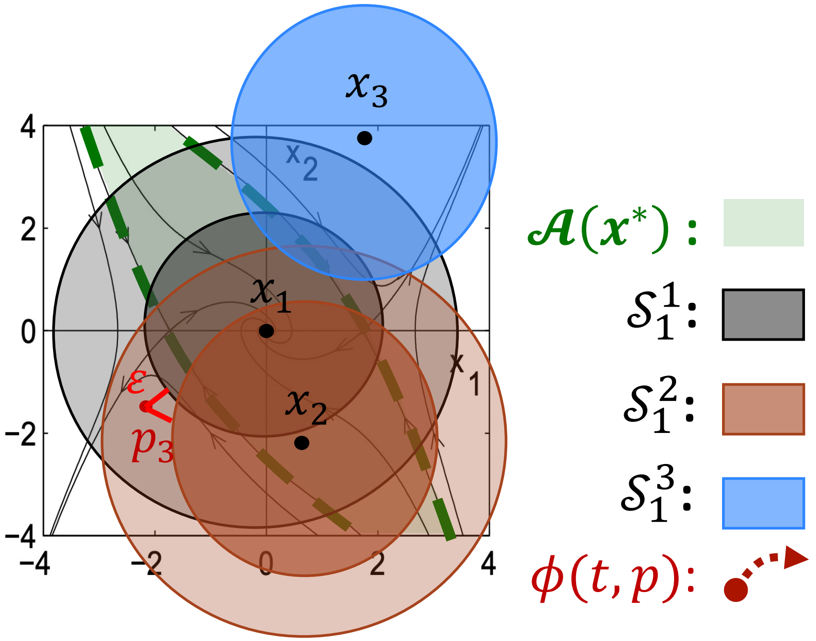

[arXiv 22] Shen, Bichuch, M, *Model-free Learning of Regions of Attraction via Recurrent Sets,* **CDC 2022,** journal preprint arXiv:2204.10372.

Example: Progressively Expanding the RoA Approximation

- At Each Episode:
  - **Sample 50** center points (uniformly)
  - **Stopping criteria:** $\delta = 10^{-5}$

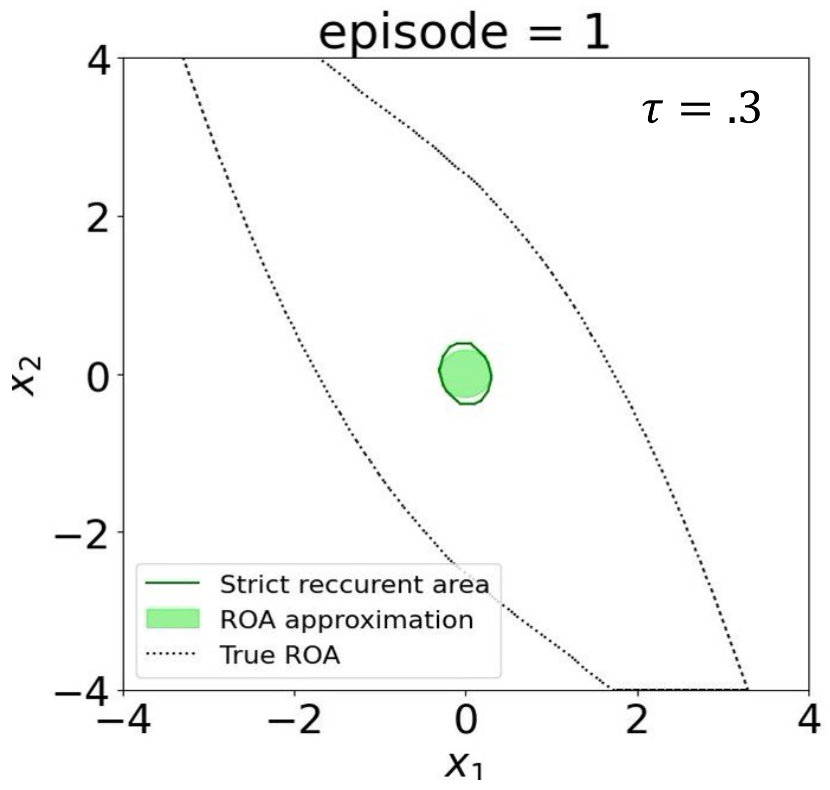

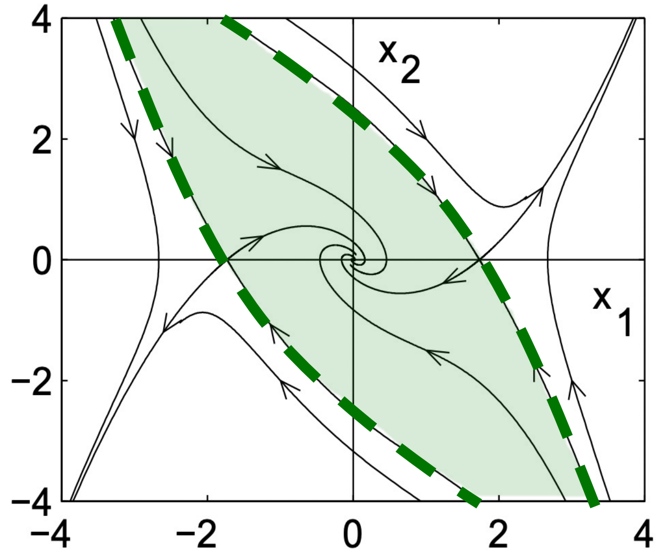

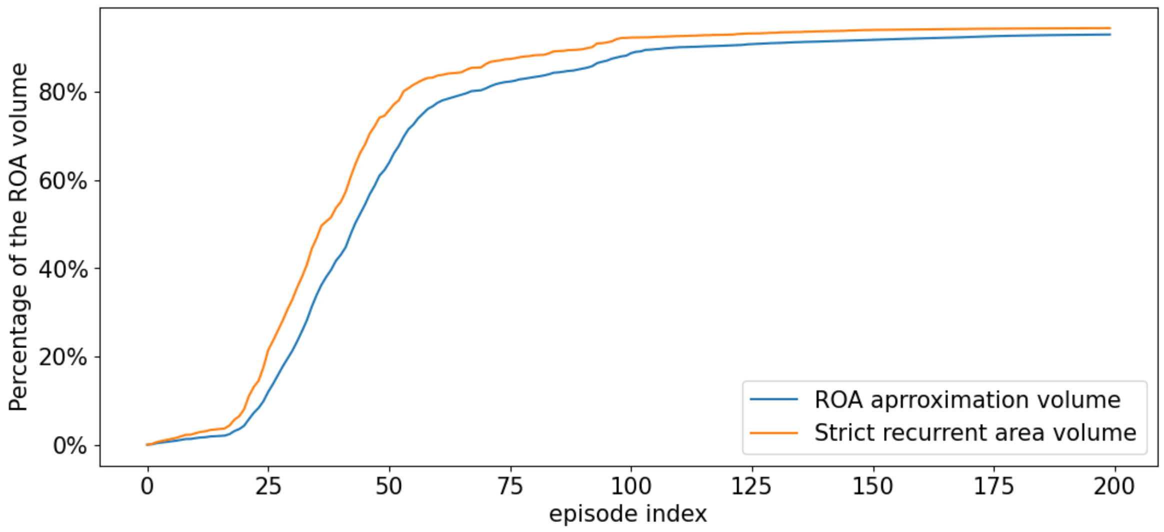

Outline

- Invariance: Merits and trade-offs
- Letting things go, and come back: Recurrent sets
- **Analysis using recurrent sets**
  - Approximating regions of attractions
  - **Stability analysis via non-monotonic Lyapunov functions**
- Recurrence in nonlinear control systems
  - Entropy and bit rates of control recurrent sets

Lyapunov's Direct Method

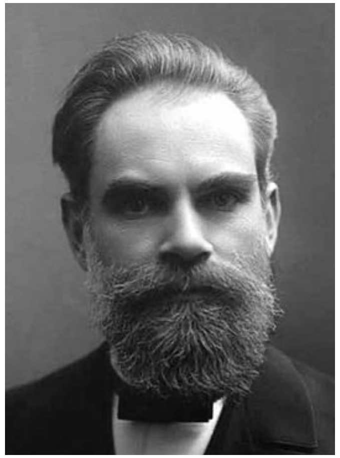

Key idea: Make sub-level sets invariant to trap trajectories

> ***Theorem [Lyapunov '1892]****. Given $V:\mathbb{R}^d \to \mathbb{R}_{\geq 0}$, with $V(x) > 0, \forall x \in \mathbb{R}^d \backslash \{x^*\}$, then:*
> - $\dot{V} \leq 0 \to x^*$ *stable*
> - $\dot{V} < 0 \to x^*$ *as. stable*

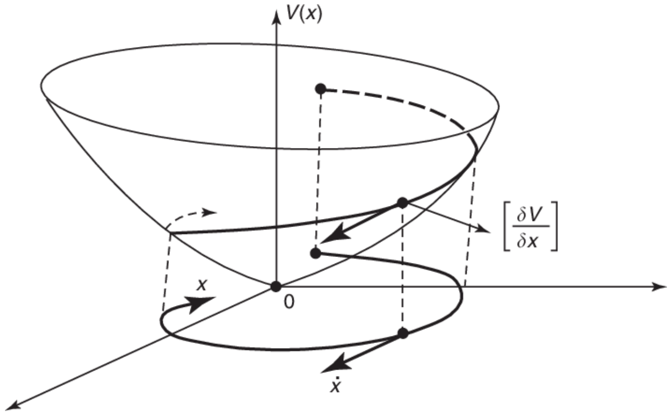

Challenge: Couples shape of V and vector field f

- Towards decoupling the $V - f$ geometry
  - Controlling regions where $\dot{V} \geq 0$ **[Karafyllis '09, Liu et al '20]**
  - Higher order conditions: $g(V^{(q)}, \dots, \dot{V}, V) \leq 0$ **[Butz '69, Gunderson '71, Ahmadi '06, Meigoli '12]**
  - Discretization approach: $V\big(x(T)\big) \leq V(x(0))$ **[Coron et al '94, Aeyels et. al '98, Karafyllis '12]**
  - Multiple Lyapunov Functions: $\{V_j : j \in [k]\}$ **[Ahmadi et al '14]**

A Butz. Higher order derivatives of Lyapunov functions. IEEE Transactions on automatic control, 1969
Gunderson. A comparison lemma for higher order trajectory derivatives. Proceedings of the American Mathematical Society, 1971
Coron, Lionel Rosier. A relation between continuous time-varying and discontinuous feedback stabilization. J. Math. Syst., Estimation, Control, 1994
Aeyels, Peuteman. A new asymptotic stability criterion for nonlinear time-variant differential equations. IEEE Transactions on automatic control, 1998
Ahmadi. Non-monotonic Lyapunov functions for stability of nonlinear and switched systems: theory and computation, 2008
Karafyllis, Kravaris, Kalogerakis. Relaxed Lyapunov criteria for robust global stabilisation of non-linear systems. International Journal of Control, 2009
Meigoli, Nikravesh. Stability analysis of nonlinear systems using higher order derivatives of Lyapunov function candidates. Systems & Control Letters, 2012
Karafyllis. Can we prove stability by using a positive definite function with non sign-definite derivative? IMA Journal of Mathematical Control and Information, 2012
Ahmadi, Jungers, Parrilo, Roozbehani. Joint spectral radius and path-complete graph Lyapunov functions. SIAM Journal on Control and Optimization, 2014
Liu, Liberzon, Zharnitsky. Almost Lyapunov functions for nonlinear systems. Automatica, 2020

Lyapunov's Direct Method

Key idea: Make sub-level sets invariant to trap trajectories

Theorem [Lyapunov '1892]. *Given* $V: \mathbb{R}^d \to \mathbb{R}_{\geq 0}$, *with* $V(x) > 0, \forall x \in \mathbb{R}^d \backslash \{x^*\}$, *then:*

- $\dot{V} \leq 0 \to x^*$ *stable*
- $\dot{V} < 0 \to x^*$ *as. stable*

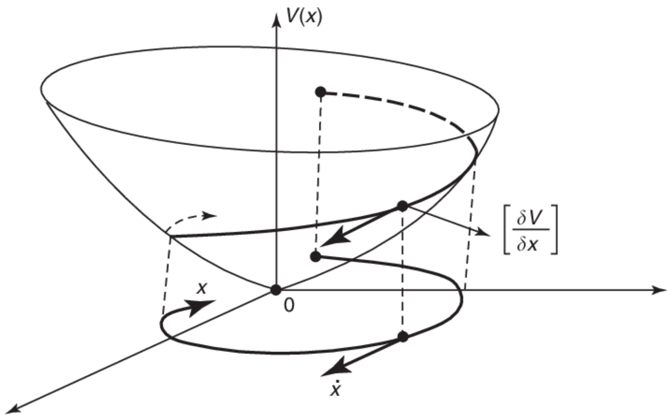

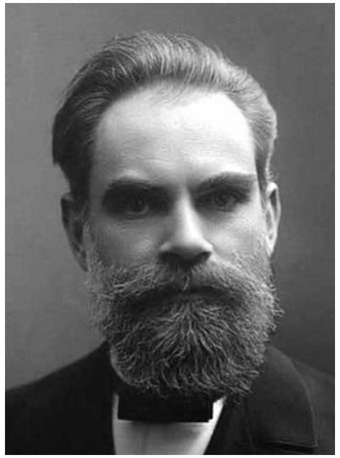

Challenge: Couples shape of V and vector field f

- Towards decoupling the $V - f$ geometry
  - Controlling regions where $\dot{V} \geq 0$ **[Karafyllis '09, Liu et al '20]**
  - Higher order conditions: $g(V^{(q)}, \ldots, \dot{V}, V) \leq 0$ **[Butz '69, Gunderson '71, Ahmadi '06, Meigoli '12]**
  - Discretization approach: $V\big(x(T)\big) \leq V(x(0))$ **[Coron et al '94, Aeyels et. al '98, Karafyllis '12]**
  - Multiple Lyapunov Functions: $\{V_j : j \in [k]\}$ **[Ahmadi et al '14]**

Question: Can we provide stability conditions based on recurrence?

Recurrently Decreasing Lyapunov Functions

A continuous function $V: \mathbb{R}^d \to \mathbb{R}_+$ is a **recurrently non-increasing Lyapunov function** over intervals of length $\boldsymbol{\tau}$ if

$$\mathcal{L}_f^{(0,\tau]} V(x) := \min_{t\in(0,\tau]} V(\phi(t,x)) - V(x) \leq 0 \quad \forall x \in \mathbb{R}^d$$

Preliminaries:

- Sub-level sets $\{V(x) \leq c\}$ are $\boldsymbol{\tau}$-recurrent sets.
- When f is L-Lipschitz, one can trap trajectories.

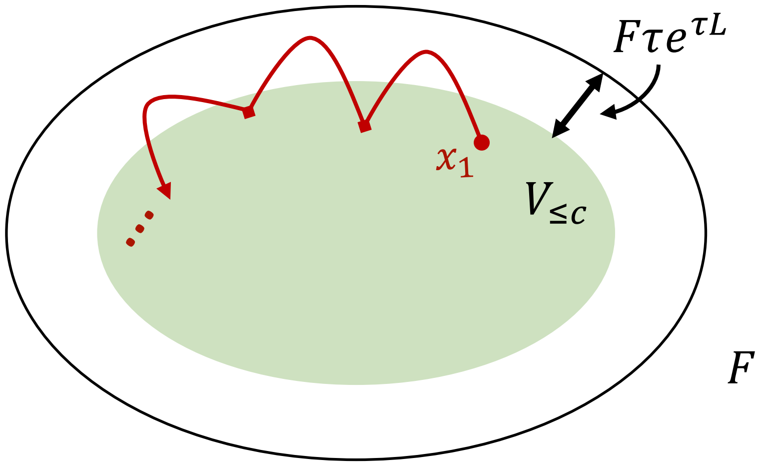

$F = \max_{x\in S} ||f(x)||$

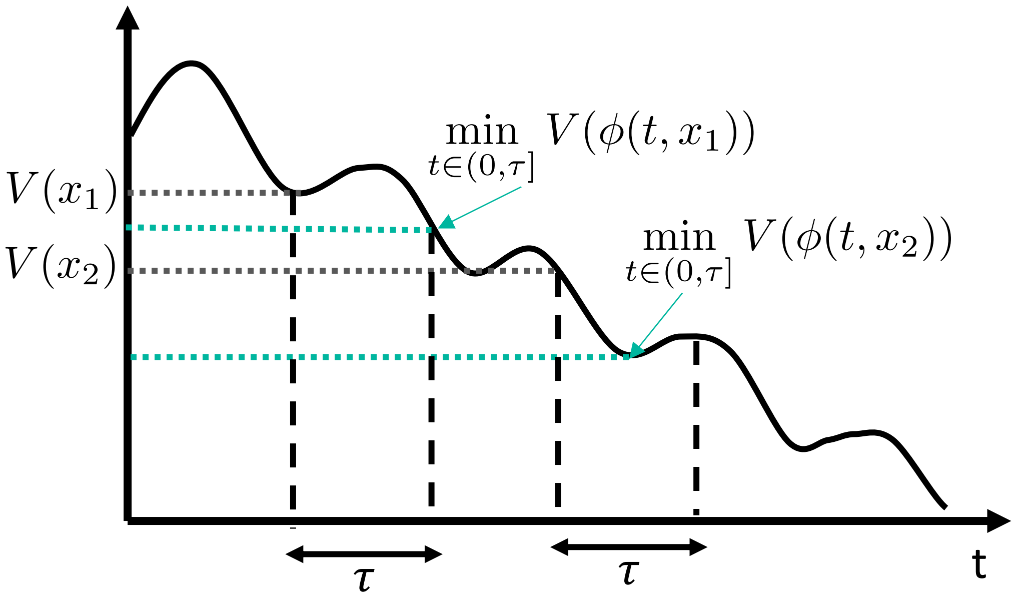

Recurrently Non-Increasing Lyapunov Functions

A continuous function $V\colon \mathbb{R}^d \to \mathbb{R}_+$ is a **recurrently non-increasing Lyapunov function** over intervals of length $\boldsymbol{\tau}$ if

$$\mathcal{L}_f^{(0,\tau]} V(x) := \min_{t\in(0,\tau]} V(\phi(t,x)) - V(x) \le 0 \quad \forall x \in \mathbb{R}^d \qquad \textbf{(*)}$$

Theorem [CDC 23*]: *Let $V\colon \mathbb{R}^d \to \mathbb{R}_{\ge 0}$ be a recurrently non-increasing Lyapunov function over intervals of length τ. Let f be L-Lipschitz*

- *Then the equilibrium x^* is stable.*
- *Further, if the* ***inequality is strict****, then x^* is asymptotically stable!*

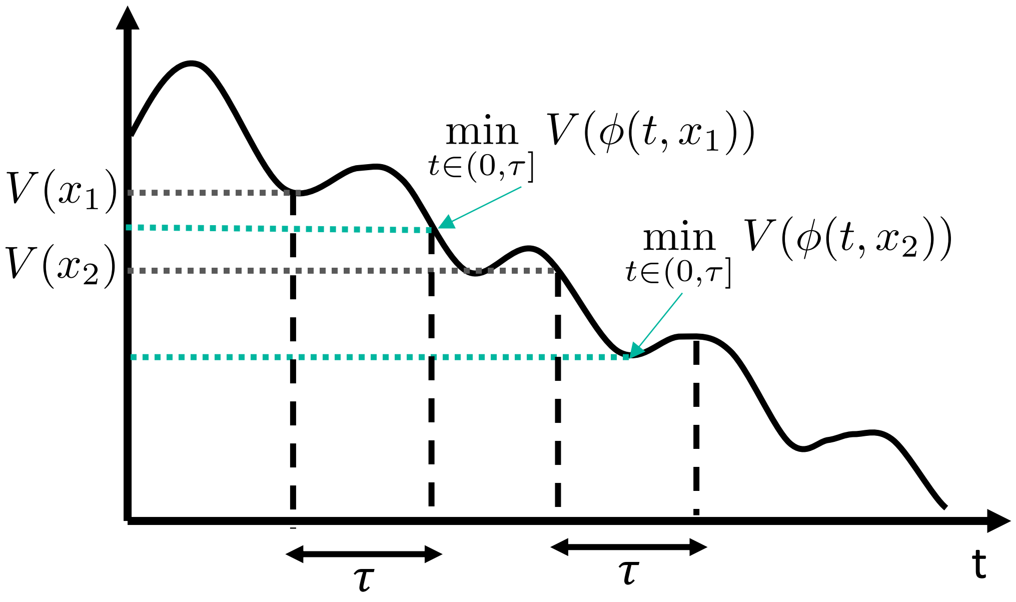

Siegelmann, Shen, Paganini, M, *A recurrence-based direct method for stability analysis and GPU-based verification of non-monotonic Lyapunov functions,* **submitted CDC 2023**

Exponential Stability Analysis

The function $V: \mathbb{R}^d \to \mathbb{R}_+$ is ***α*-exponentially recurrently *τ*-decreasing Lyapunov function** over intervals of length $\boldsymbol{\tau}$ if

$$\mathcal{L}_{f,\alpha}^{(0,\tau]} V(x) := \min_{t \in (0,\tau]} e^{\alpha t} V(\phi(t,x)) - V(x) \leq 0 \quad \forall x \in \mathbb{R}^d$$

Theorem [CDC 23*]: *Let* $V: \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ *satisfy*

$$\alpha_1 ||x - x^*|| \leq V(x) \leq \alpha_2 ||x - x^*||.$$

Then, if V is α-exponentially recurrently τ-decreasing Lyapunov function, then x^* is exponentially stable with rate α.

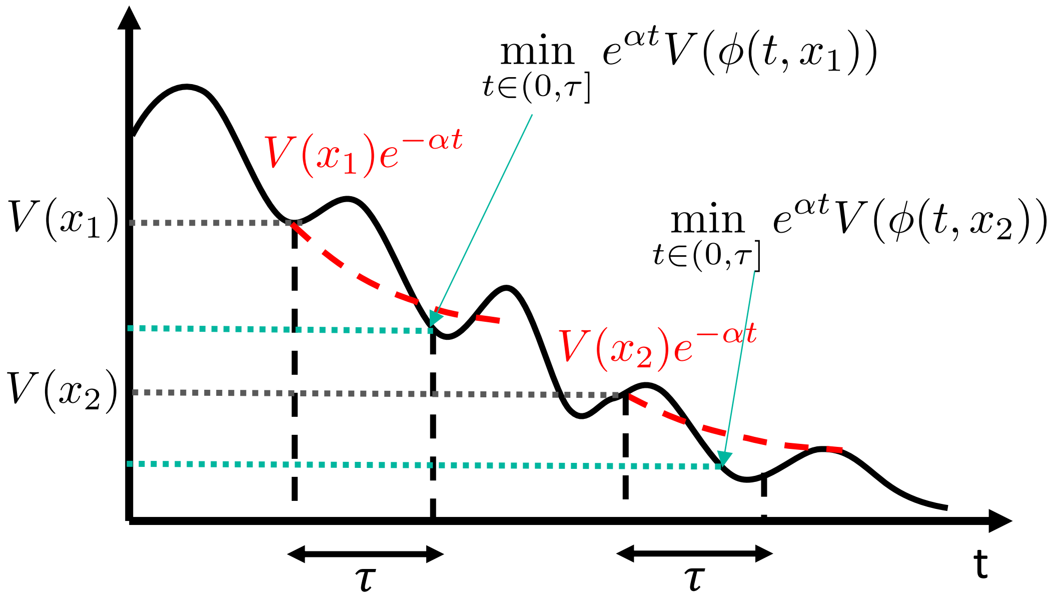

Siegelmann, Shen, Paganini, M, *A recurrence-based direct method for stability analysis and GPU-based verification of non-monotonic Lyapunov functions,*
***submitted CDC 2023**

All norms are Lyapunov functions!

Theorem: Assume x^* is globally exponentially stable: $\exists\, K, c > 0$ such that:

$$\left|\left|\phi(t,x) - x^*\right|\right| \leq K e^{-ct} \left|\left|x_0 - x^*\right|\right|.$$

Then, $V(x) = \left|\left|x - x^*\right|\right|$ *is* *α-exponentially recurrently τ-decreasing*, *i.e.,*

$$\min_{t\in(0,\tau]} e^{\alpha t}\left|\left|\phi(t,x) - x^*\right|\right| - \left|\left|x - x^*\right|\right| \leq 0, \qquad \forall x \in \mathbb{R}^d,$$

whenever $\alpha < c$ *and* $\tau \geq \frac{1}{c-\alpha}\ln K.$

Remarks:

- The rate α must be strictly smaller than the rate of convergence c (giving up optimality).
- *Any norm is a Lyapunov function!*

Question: Is the struggle for its search over?

Outline

- Invariance: Merits and trade-offs
- Letting things go, and come back: Recurrent sets
- Analysis using recurrent sets
  - Approximating regions of attractions
  - Stability analysis via non-monotonic Lyapunov functions
- **Recurrence in nonlinear control systems**
  - Entropy and bit rates of control recurrent sets

Various notions of entropy in the literature

Mainly bounding the bit rates needed to perform various estimation and control tasks over limited-bandwidth channels.

Examples:

- Topological entropy **[Adler 1965, Bowen 1971, Savkin 2006]**
- Estimation entropy **[Liberzon and Mitra 2016, 2018, Sibai and Mitra 2017, 2018, 2023]**
- Stabilization entropy **[Colonius 2012, Nair et al. 2004]**
- Invariance entropy **[Colonius and Kawan 2009, 2011, Rungger and Zamani 2017, Tomar et al. 2021, 2022]**

Controlled recurrent sets: Letting things go, and come back

Problem Setup:

- Continuous time controlled dynamical system: $\dot{x}(t) = f(x(t), u(t))$
- Initial condition $x_0 = x(0)$, solution at time t: $\phi(t, x_0, u)$.

A set $\mathcal{R} \subseteq \mathbb{R}^d$ is **controlled τ-recurrent**, for some $\tau \geq 0$, if for any $x_0 \in \mathcal{R}$, $\exists u \in \mathcal{U}$, $\exists t \in (0, \tau]$ s.t. $\phi(t, x_0, u) \in \mathcal{R}$.

Recurrent trajectories: they go, and come back(ish)

Similarly to other entropy notions, we require a relaxed notion of recurrence…

Definition: $(\boldsymbol{T}, \boldsymbol{\varepsilon}, \boldsymbol{\tau}, \boldsymbol{\mathcal{R}})$-recurrence
Fix any $\tau \geq 0$, $\varepsilon \geq 0$, $T \geq \tau$, $x_0 \in \mathcal{R}$, and $u \in \mathcal{U}$. The trajectory ξ is **$(\boldsymbol{T}, \boldsymbol{\varepsilon}, \boldsymbol{\tau}, \boldsymbol{\mathcal{R}})$-recurrent**, if $\forall t \in [0, T-\tau]$, $\exists t' \in [t, t+\tau]$ such that $\xi(t', x, u) \in B_\varepsilon(\mathcal{R})$.

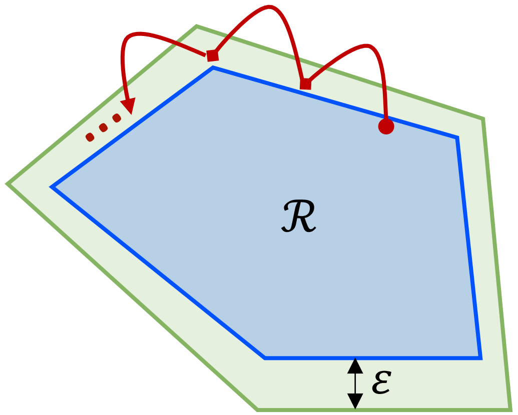

Bloated recurrent set $B_\varepsilon(\mathcal{R})$:

Controlled Recurrent set $\mathcal{R}$:

A recurrent trajectory:

Spanning sets

We define open-loop control signals sufficient for (almost) recurrence

Definition: $(\boldsymbol{T}, \boldsymbol{\varepsilon}, \boldsymbol{\tau}, \boldsymbol{\mathcal{R}})$-spanning Set
A set $S \subseteq \mathcal{U}$ is called a recurrence $(\boldsymbol{T}, \boldsymbol{\varepsilon}, \boldsymbol{\tau}, \boldsymbol{\mathcal{R}})$**-spanning** set if for any $x_0 \in \mathcal{R}$, there exists a $u \in S$ such that ξ is $(\boldsymbol{T}, \boldsymbol{\varepsilon}, \boldsymbol{\tau}, \boldsymbol{\mathcal{R}})$-recurrent.

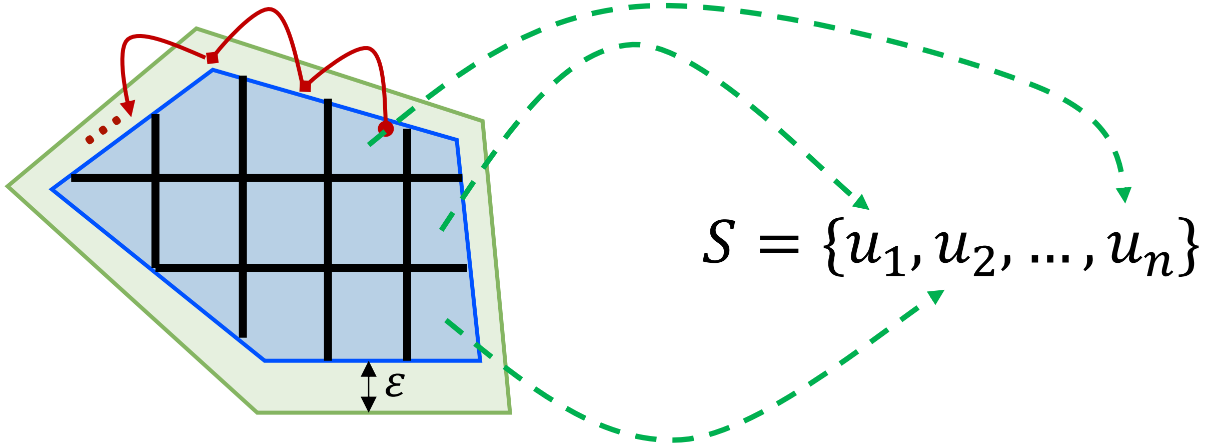

Bloated recurrent set $B_\varepsilon(\mathcal{R})$:

Controlled Recurrent set $\mathcal{R}$:

A recurrent trajectory:

Mapping states to control signals:

Set of states mapped to the same control signal:

Recurrence entropy

Definition: Recurrence entropy

$$h_{\text{rec}}(\tau,\mathcal{R}) := \lim_{\varepsilon \searrow 0} \limsup_{T\to\infty} \frac{1}{T} \log r_{\text{rec}}(T,\varepsilon,\tau,\mathcal{R}),$$

where $r_{\text{rec}}(T,\varepsilon,\tau,\mathcal{R})$ is the minimal cardinality of a spanning set.

Remark: Measures the exponential rate at which the number of (open-loop) control signals needed to achieve recurrence increases as time horizon T and recurrence strictness ε^{-1} increase.

Relation to Invariance Entropy

Existing notion of invariance entropy, i.e., $h_{\mathrm{inv}}(X_0, \mathcal{R})$, where $X_0 \subseteq \mathcal{R}$, is a special case of recurrence entropy

Proposition:

$$h_{\mathrm{inv}}(X_0, \mathcal{R}) := h_{\mathrm{rec}}(0, \mathcal{R}) = \lim_{\varepsilon \searrow 0} \limsup_{T \to \infty} \frac{1}{T} \log r_{\mathrm{rec}}(T, \varepsilon, 0, \mathcal{R}),$$

where $r_{\mathrm{rec}}(T, \varepsilon, 0, \mathcal{R})$ is the minimal cardinality of a spanning set that keeps $\mathrm{B}_{\varepsilon}(\mathcal{R})$ **invariant**, i.e., recurrent with $\tau = 0$.

Questions:

- How different are $h_{\mathrm{inv}}(\mathcal{R})$ and $h_{\mathrm{rec}}(\tau, \mathcal{R})$?
- How does $h_{\mathrm{rec}}(\tau, \mathcal{R})$ change as τ increases?

[1] Colonius, Kawan. Invariance entropy for control systems. SIAM Journal on Control and Optimization, 2009

Relation between Invariance and Recurrence

Theorem: Assume $\mathcal{R}$ is controlled invariant, then:

$$h_{\text{inv}}\big(\mathcal{R}, B_{\delta_\tau}(\mathcal{R})\big) \leq h_{\text{rec}}\;(\tau, \mathcal{R}) \leq h_{\text{inv}}(\mathcal{R}, \mathcal{R})$$

where $\delta_\tau = \tau e^{L_\tau \tau} F_{\mathcal{R}}$ is a constant dependent on τ, f, and $\mathcal{R}$ and L_τ is a locally Lipschitz constant of the vector field f.

Proof: (sketch)

- *Left inequality:* containment lemma (bounding distance from recurrent trajectories to $\mathcal{R}$)
- *Right inequality:* any invariance causing control is also recurrence enforcing.

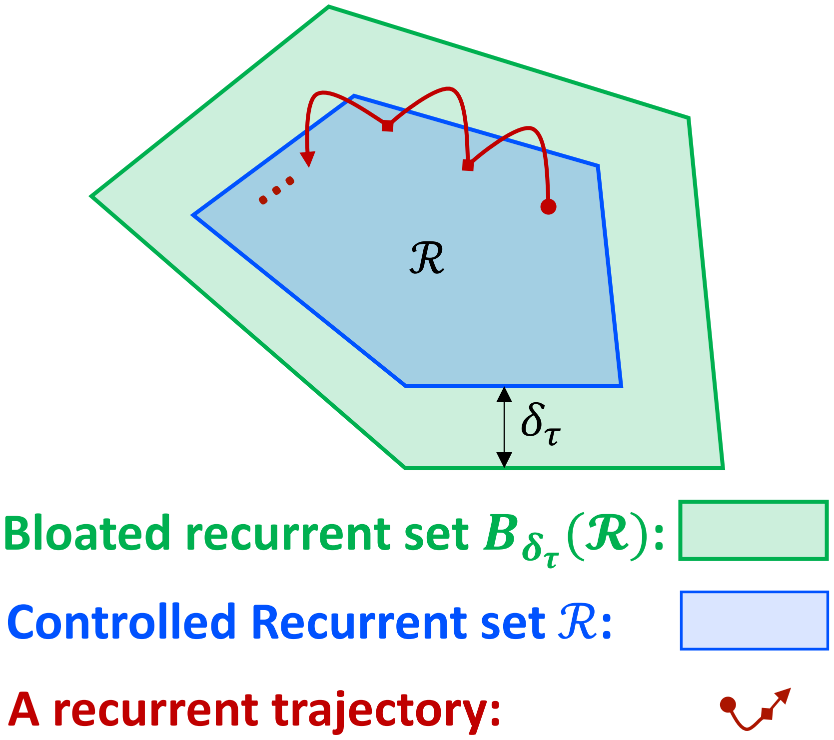

Example of strict separation between $h_{\text{inv}}(\mathcal{R},\mathcal{R})$ and $h_{\text{rec}}(\tau,\mathcal{R})$

Consider the system:

$$\begin{bmatrix}\dot{x}_1\\ \dot{x}_2\end{bmatrix} = \begin{bmatrix}0 & 1\\ 0 & 0\end{bmatrix}\begin{bmatrix}x_1\\ x_2\end{bmatrix} + \begin{bmatrix}0\\ 1\end{bmatrix} u$$

where $u \in U = [-1,1]$.

Now, consider the controlled recurrent set $\mathcal{R} = [-1,1]^2$.

Theorem: $h_{\text{inv}}(\mathcal{R},\mathcal{R}) = \infty$ and $h_{\text{rec}}(\tau,\mathcal{R}) \begin{cases} = \infty, & \tau < 2 \\ \leq \frac{2}{\ln 2}, & \tau \geq 0 \end{cases}$

Bound on Recurrence Entropy

Theorem: Bounds on $h_{\text{rec}}(\tau,\mathcal{R})$
Whenever $\mathcal{R}$ is a controlled τ-recurrent set. Then for any $\tau' \geq \tau$:

$$\frac{1}{\ln 2}\left[\min_{(x,u)\in B_{\delta_{\tau'}}(\mathcal{R})\times U} \operatorname{div}_x f(x,u)\right]_+ \leq h_{\text{rec}}(\tau',\mathcal{R}) \leq h_{\text{rec}}(\tau,\mathcal{R}) \leq \frac{L_\tau n}{\ln 2}$$

Remarks:

- When $\tau = \tau' = 0$, we recover the bounds on invariance entropy by Colonius and Kawan 2012.
- If a set is controlled τ-recurrent, making the set τ'-recurrent is at most as hard as making it τ-recurrent.
- Moreover, as $\tau' \to \infty$, the lower bound goes to zero, as expected.

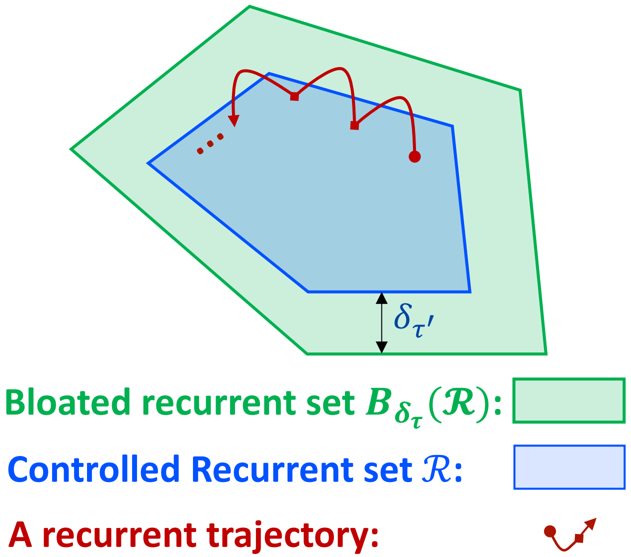

Bit rates needed to enforce recurrence

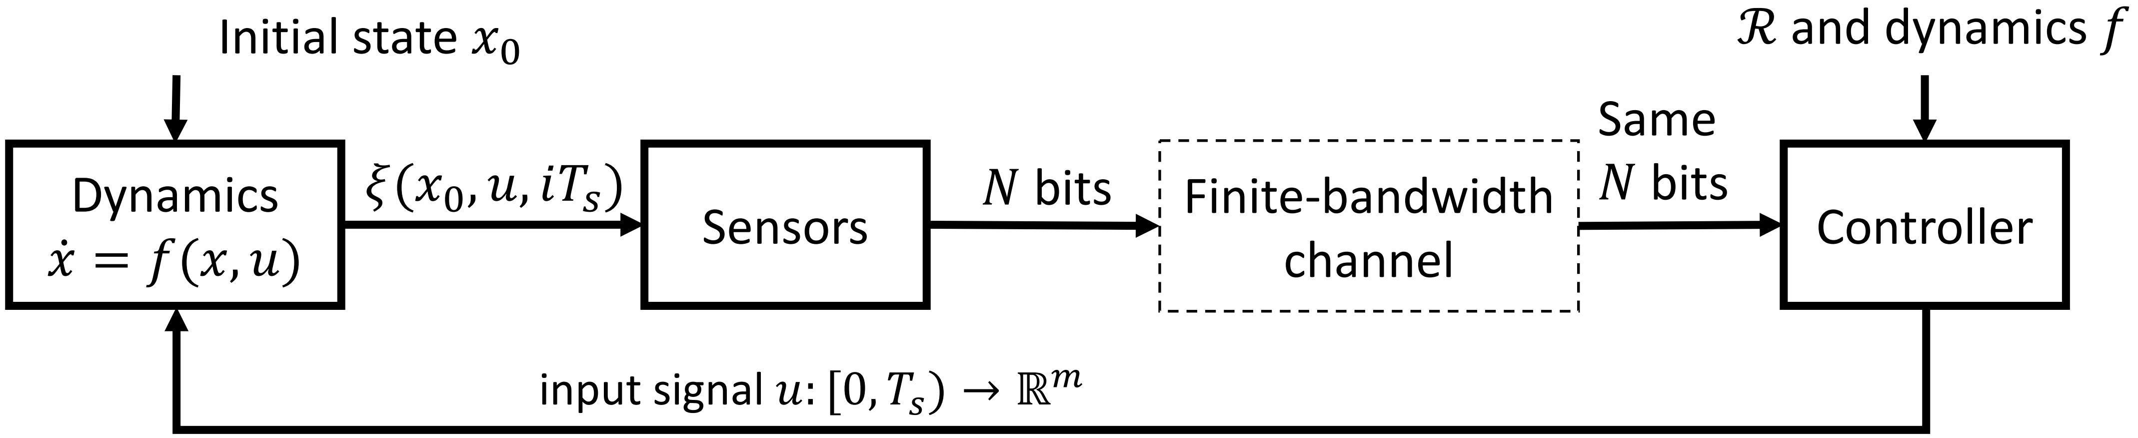

Problem: Given ε ∈ $\mathbb{R}^{>0}$, what is the **minimum bit rate** N/T_s needed for $\xi(x_0, u, t)$ to be $(\boldsymbol{\varepsilon}, \boldsymbol{\tau}, \boldsymbol{\mathcal{R}})$-recurrent?

Theorem: For any $\varepsilon \geq 0$, there exists no $(\varepsilon, \tau, \mathcal{R})$-recurrence enforcing algorithm with an average bit rate smaller than $h_{\text{rec}}(\tau, \mathcal{R})$.

Algorithm

Enforcing (asymptotic) τ-recurrence over limited-bandwidth channels

Algorithm 1 Sensor algorithm for achieving recurrence

input: Q, $\varepsilon \in (0, \varepsilon^*]$, $\tau > 0$, $g : B_{\delta_\tau+\varepsilon}(Q) \times \mathbb{R}^{\geq 0} \to U$
$S_0 \leftarrow Q$
$r_0 \leftarrow \varepsilon$
$C_0 \leftarrow grid(S_0, r_0 e^{-(L_\tau+\alpha)\tau})$
$i = 0$
while true **do**
    $x_i \leftarrow sense()$
    $q_i \leftarrow quantize(x_i, C_i)$
    $send(encode(q_i, C_i))$
    $u_i \leftarrow g(q_i, [0, \tau))$
    $r_{i+1} \leftarrow r_i e^{-\alpha\tau}$
    $S_{i+1} \leftarrow B_{r_{i+1}}(simulate(q_i, u_i, \tau))$
    $C_{i+1} \leftarrow grid(S_{i+1}, r_{i+1} e^{-(L_\tau+\alpha)\tau})$
    $i \leftarrow i + 1$
    $sleep(\tau)$

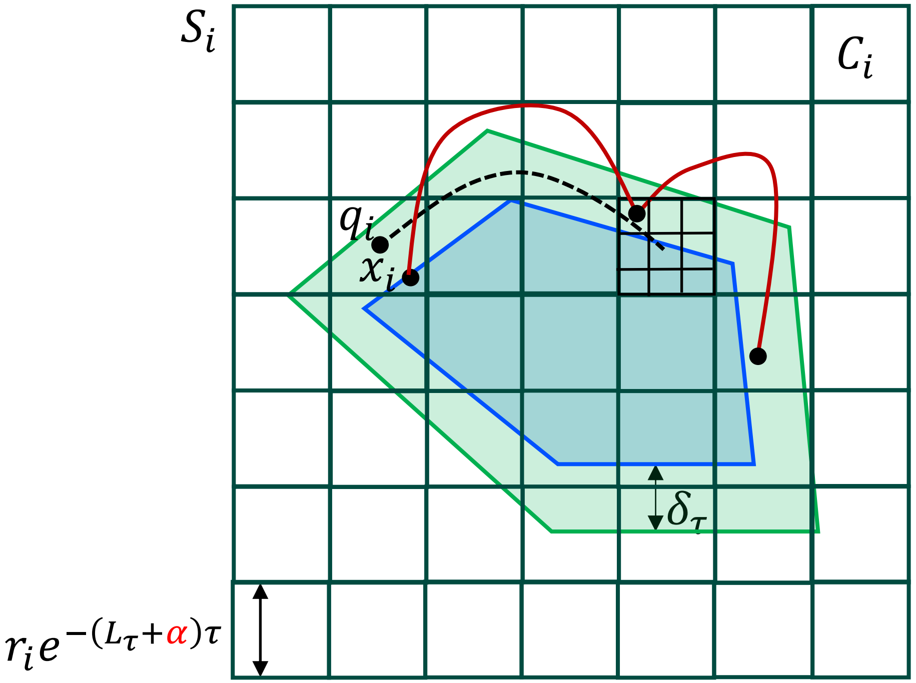

Theorem: Algorithm 1 guarantees that starting from any state $x_0 \in \mathcal{R}$, the trajectory of the system will converge to a (τ, Q)-recurrent trajectory at an exponential rate of α. It requires an average bit rate of $\frac{n(L_\tau+\alpha)}{\ln 2}$ between the sensor and the actuator.

Conclusions and Future work

- **Takeaways**
  - Proposed a **relaxed notion of invariance** known as **recurrence.**
  - Provide **necessary and sufficient conditions** for a recurrent set to be an **inner approximation** of the RoA.
  - Generalized Lyapunov Theory **for recurrently decreasing functions** using recurrent sets
  - From an information theoretical standpoint, **making as set recurrent can be easier than invariant.**

- **Ongoing work**
  - **Recurrent Sets:** Smart choice of multi-points, control recurrent sets, GPU implementation
  - **Lyapunov Functions:** Generalize other Lyapunov notions, Control Lyapunov Functions, Barrier Functions, Control Barrier Functions, Contraction, etc.
  - **Entropy:** Understanding the memory complexity of making a set recurrent and generalizations to other tasks

Thanks!

Related Publications:

[arXiv 22] Shen, Bichuch, M, *Model-free Learning of Regions of Attraction via Recurrent Sets,* **CDC 2022, journal preprint arXiv:2204.10372**.
[CDC 23] Siegelmann, Shen, Paganini, M, *A recurrence-based direct method for stability analysis and GPU-based verification of non-monotonic Lyapunov functions*, **CDC 2023**
[HSCC 24] Sibai, M, *Recurrence of nonlinear control systems: Entropy and bit rates,* **HSCC, 2024**

Enrique Mallada
mallada@jhu.edu
http://mallada.ece.jhu.edu